



Regularising experimental correlations in LHC data: theory and application to a global analysis of parton distributions

Zahari Kassabov¹, Emanuele R. Nocera^{2,a} , Michael Wilson²

¹ DAMTP, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, UK

² The Higgs Centre for Theoretical Physics, University of Edinburgh JCMB, Mayfield Rd, Edinburgh, KB EH9 3JZ, UK

Received: 8 July 2022 / Accepted: 17 October 2022 / Published online: 27 October 2022
© The Author(s) 2022

Abstract We show how an inaccurate determination of experimental uncertainty correlations in high-precision LHC measurements may undermine the reliability of the associated χ^2 . We formulate the problem rigorously, and devise a regularisation procedure that increases the stability of the χ^2 by altering the covariance matrix of the measurement as little as possible. We apply the procedure to the NNPDF4.0 global analysis of parton distribution functions that utilises a large amount of LHC measurements. We find that the regularised χ^2 of the NNPDF4.0 determination is lowered by about 3σ , without significantly altering the resulting PDFs upon refitting.

1 Introduction

With the restart of the Large Hadron Collider (LHC), particle physics is getting into the thick of a new era, whereby measurements are anticipated to attain an unprecedented, percent-level, statistical precision [1]. These measurements are utilised to improve the determination of Standard Model (SM) parameters [2], to constrain Parton Distribution Functions (PDFs) [3], to evaluate backgrounds for missing energy searches [4], or more generally to constrain higher-dimensional operators in the SM Lagrangian [5], and eventually to stress-test the properties of the Higgs boson [6].

In all of these cases, measurements are contrasted with theoretical predictions by means of statistical inference: a model is chosen and compared or optimised to the data through maximum likelihood estimation. The goal is to reject a test hypothesis or to obtain a confidence interval of model parameters. Because experimental uncertainties in LHC measurements are commonly assumed to be Gaussian, the figure of merit utilised for the statistical test or to optimise the model

is the χ^2 statistic, which is monotonic in the likelihood of sampling the experimental data given the theory.

The robustness of the χ^2 as a figure of merit relies on the accuracy of theoretical expectations and of experimental uncertainties. In this paper we assume that the χ^2 is not spoiled by inaccuracies in theoretical expectations — a fact that is possibly not true now, but that will become increasingly realistic in the future [7] — and focus only on inaccuracies of experimental uncertainties.

A proper estimation of uncertainties in LHC measurements is indeed becoming increasingly delicate. The large event samples collected during Run I and II have been making statistical uncertainties generally smaller than systematic uncertainties; the upcoming Run III will make the former even smaller. In contrast to statistical uncertainties, systematic uncertainties (which are not related to event counts but, e.g., to limitations of the detector or to assumptions made in their modelling) are more difficult to estimate. The reason being that custom procedures and subjective choices are involved in their quantification [8]. Furthermore, systematic uncertainties are usually correlated across different kinematic bins, both within the same measurement and across different measurements. Determining these correlations is even more difficult, and often it is not even attempted. In this case, simple assumptions, such as taking systematic uncertainties to be fully correlated or fully uncorrelated, may misrepresent the truth. More elaborated guesswork can be performed in order to devise *ad-hoc* correlation models, which however have no generality and can be time consuming.

Because the uncertainties on LHC measurements are being increasingly dominated by systematic uncertainties, any analysis that utilises them is implicitly dependent on the choices made in their characterisation. While this dependence is generally unavoidable, some care must be taken to prevent it from hampering the use of the data in precision physics analyses based on statistical inference.

^a e-mail: emanuele.roberto.nocera@gmail.com (corresponding author)

The aim of this paper is to formulate and address this problem rigorously. We first demonstrate how inaccuracies in the estimation of systematic uncertainty correlations, even if small, can lead to instabilities in the experimental covariance matrix and how these can ultimately undermine the reliability of the χ^2 . We then devise a regularisation procedure whereby these instabilities are removed with minimal information on their source, and without loss of generality. The idea is to define a bound on the singular values of the correlated part of the matrix of uncertainties, and to clip them to a suitably chosen value that alters only the small subset of directions associated to instability. We finally apply this procedure to a particular problem relevant to LHC precision physics that utilises statistical inference: PDF determination. Although we orient our discussion towards this problem, our regularisation procedure is completely general, and can be applied to any statistical analysis that involves the evaluation of the χ^2 .

The structure of the paper is as follows. In Sect. 2 we introduce the matrix of uncertainties and formulate a stability criterion for it. In Sect. 3 we derive our regularisation procedure and demonstrate how it works with a toy model. In Sect. 4 we apply the procedure to PDF determination using the recently released NNPDF4.0 parton sets [9]. We summarise our results in Sect. 5. The paper is completed by two appendices: Appendix A is a glossary of some useful definitions used through the paper; Appendix B contains the proof of Eq. (14). Our regularisation procedure is made publicly available as part of the NNPDF software [10].

2 Matrix of uncertainties and its stability

In this section we formulate the problem of the reliability of the χ^2 if instabilities, even if small, appear in the covariance matrix that enters its computation. We first introduce the matrix of uncertainties and write the χ^2 in terms of it. We then derive an upper bound on the instability of the matrix of uncertainties that ensures the stability of the χ^2 with minimal information.

2.1 Matrix of uncertainties and χ^2

The format of LHC measurements, as often made public through the HEPDATA repository [11], consists of a central value and of a set of uncertainties for each of the data points that form the measurement itself. The set of uncertainties usually encompass a total statistical uncertainty and a set of independent systematic uncertainties. The latter are typically correlated across data points, by an amount that may be specified or not.

Let us consider an experimental measurement made of N_{dat} data points, each of which has N_{err} independent uncer-

tainties. We call \mathbf{d} the vector of experimental mean values, $\mathbf{d} = \{D_i\}$, and A the $N_{\text{dat}} \times N_{\text{err}}$ matrix of uncertainties, $A = \{A_{ij}\}$, with $i = 1, \dots, N_{\text{dat}}$, and $j = 1, \dots, N_{\text{err}}$. Assuming that all uncertainties are Gaussian and that they are combined additively, the experimental measurement defines a multi-Gaussian distribution with mean \mathbf{d} , given by the experimental central values, and covariance matrix C , given by the product of the matrix of uncertainties and its transpose, $C = AA^t$.

Depending on the information provided with a given experimental measurement, each element of the matrix of uncertainties can be obtained from knowledge of C , for example by taking its Cholesky decomposition, or from direct knowledge of experimental uncertainties. In this latter case, should \mathcal{O}_i be the physical observable corresponding to the data points D_i , and $\{u_j\}$ the set of independent variables which contribute to the experimental uncertainty and on which the observable depends (each described by a Gaussian distribution with central value u_j^0 and uncertainty s_j), any element of the matrix of uncertainties reads as

$$A_{ij} = \left. \frac{\partial \mathcal{O}_i}{\partial u_j} \right|_{\mathbf{u}=\mathbf{u}^0} s_j. \quad (1)$$

If a given source of uncertainty u_l affects a single data point k , then $\partial \mathcal{O}_i / \partial u_l = 0$ for $i \neq k$, and it corresponds to a row in A with a single non-zero entry A_{kl} . For instance, this is the case for statistical uncertainties that originate from bin-by-bin event counts. These uncertainties, together with similarly fully uncorrelated systematic uncertainties, can therefore be encoded in a $N_{\text{dat}} \times N_{\text{dat}}$ diagonal sub-matrix of A . We assume that such uncertainties are always present in a measurement, therefore we will henceforth consider that $N_{\text{err}} \geq N_{\text{dat}}$, and that both A and C be full rank.

The inverse of the covariance matrix C is $C^{-1} = A^+ A^+$, where A^+ is the right inverse of A (see Appendix A). Denoting with $\mathbf{t} = \{T_i\}$, $i = 1, \dots, N_{\text{dat}}$, the vector of theoretical predictions corresponding to the data mean values \mathbf{d} , the χ^2 can be written as

$$\chi^2 = (\mathbf{d} - \mathbf{t})^t C^{-1} (\mathbf{d} - \mathbf{t}) = \|A^+ (\mathbf{d} - \mathbf{t})\|^2. \quad (2)$$

In this equation we have explicitly factorised the two contributions that determine the value of the χ^2 : the difference between the mean experimental central values and the theoretical expectation values, encoded in $\mathbf{d} - \mathbf{t}$; and the experimental uncertainties, encoded in the error matrix A .

Concerning the $(\mathbf{d} - \mathbf{t})$ term in Eq. (2), we assume perfect knowledge of theoretical expectations. This means that the vector of differences $\mathbf{d} - \mathbf{t}$ is a realisation of a random variable which follows a multivariate Gaussian distribution with mean zero and covariance matrix C . The corresponding probability density can be given in terms of the matrix of uncertainties

A and of a vector of N_{err} independent standard Gaussian random variables, $\mathbf{n} = \{n_j\}, j = 1, \dots, N_{\text{err}}$,

$$\mathbf{d} - \mathbf{t} = \mathbf{A}\mathbf{n}, \quad \mathbf{n} \sim \mathcal{N}(\mathbf{0}, I). \tag{3}$$

Concerning the matrix of uncertainties A in Eq. (2), we consider two different cases. The first case corresponds to assuming that A has been estimated accurately. Substituting Eq. (3) in Eq. (2), we obtain that the expected value of the χ^2 over samples of $\mathbf{d} - \mathbf{t}$ is

$$\langle \chi^2 \rangle = \langle \|A^+ \mathbf{A}\mathbf{n}\|^2 \rangle. \tag{4}$$

Using the fact that, for independent standard Gaussian variables, $\langle n_j n_l \rangle = \delta_{jl}$, we find that $\langle \chi^2 \rangle$ is given in terms of the Frobenius norm (see Appendix A) of $A^+ A$:

$$\langle \chi^2 \rangle = \sum_{j,l}^{N_{\text{err}}} (A^+ A)_{j,l}^2 = \|A^+ A\|_F^2 = N_{\text{dat}}, \tag{5}$$

where the last equality follows from the singular value decomposition of A^+ , see Appendix A.

The second case corresponds to assuming that there are inaccuracies in the estimation of uncertainties, which do not need to be large. We define as \bar{A} the matrix of uncertainties that contains such inaccuracies. This is different from A , which is therefore unknown. The covariance matrix used to compute the χ^2 is now $\bar{C} = \bar{A}\bar{A}^t$. However, because we still assume perfect knowledge of theoretical expectations, Eq. (3) continues to hold. Therefore, in analogy with Eqs. (4)–(5), the expectation value of the χ^2 reads as

$$\langle \bar{\chi}^2 \rangle = \|\bar{A}^+ A\|_F^2. \tag{6}$$

A comparison between Eqs. (6) and (5) allows one to formulate a stability criterion for the expectation value of the χ^2 and for the matrix of uncertainties A upon substituting A with \bar{A} , as we explain next.

2.2 Stability criterion

We state that the matrix of uncertainties A is stable upon the replacement $A \rightarrow \bar{A}$ in the computation of the χ^2 when the difference in its expected value, $\Delta\chi^2$, is smaller than statistical fluctuations of the χ^2 statistic itself, as measured by the standard deviation of the corresponding χ^2 distribution, which is equal to $\sqrt{2N_{\text{dat}}}$. We can therefore write a stability criterion for the expectation value of the χ^2 as:

$$\Delta\chi^2 = \left| \langle \bar{\chi}^2 \rangle - \langle \chi^2 \rangle \right| < \sqrt{2N_{\text{dat}}}. \tag{7}$$

Substituting Eqs. (5) and (6) in Eq. (7), we can equivalently write

$$\|\bar{A}^+ A\|_F^2 - N_{\text{dat}} < \sqrt{2N_{\text{dat}}}. \tag{8}$$

We now seek to find an upper bound to the inaccuracies of the matrix \bar{A} that satisfies the stability criterion on the χ^2 , Eq. (8), using minimal information. To this purpose, we write the matrix of uncertainties A , which we do not know, as a perturbation to the matrix \bar{A} , which we are given,

$$A = \bar{A} + \delta F, \tag{9}$$

where F is a matrix of perturbations and δ is a scalar parameter controlling the size of the fluctuation. We assume that δ is sufficiently small that we can linearly expand around $\delta = 0$. Replacing Eq. (9) into Eq. (6), we find

$$\langle \bar{\chi}^2 \rangle = \|\bar{A}^+(\bar{A} + \delta F)\|_F^2. \tag{10}$$

Using the triangle inequality, we can derive the upper bound

$$\langle \bar{\chi}^2 \rangle \leq (\|\bar{A}^+ \bar{A}\|_F + \delta \|\bar{A}^+ F\|_F)^2. \tag{11}$$

Then expanding the square, and using the fact that \bar{A} is full rank since it corresponds to the covariance matrix obtained in the experimental analysis, we obtain

$$\langle \bar{\chi}^2 \rangle \leq N_{\text{dat}} + 2\delta\sqrt{N_{\text{dat}}} \|\bar{A}^+ F\|_F + (\delta^2). \tag{12}$$

Now, neglecting the quadratic terms in δ , we arrive at

$$\Delta\chi^2 \leq 2\delta\sqrt{N_{\text{dat}}} \|\bar{A}^+ F\|_F. \tag{13}$$

We apply the following inequality

$$\|XY\|_F \leq \min(\|X\|_F \|Y\|_2, \|X\|_2 \|Y\|_F), \tag{14}$$

which holds for arbitrary matrices X and Y of compatible shape (see Appendix B for a proof), to Eq. (13) and find

$$\Delta\chi^2 \leq 2\delta\sqrt{N_{\text{dat}}} \|\bar{A}^+\|_2 \|F\|_F, \tag{15}$$

where $\|\bar{A}^+\|_2$ denotes the Euclidean (or L^2) norm of \bar{A} (see Appendix A). Choosing $\|\bar{A}^+\|_2 \|F\|_F$ instead of $\|\bar{A}^+\|_F \|F\|_2$ as bound in Eq. (15) results in tighter constraints. The reason being that, in practice, instabilities occur when \bar{A}^+ has large singular values.

Finally, combining Eq. (15) with Eq. (7), we conclude that the condition

$$\|\bar{A}^+\|_2 \|F\|_F \leq \frac{1}{\sqrt{2}\delta} \tag{16}$$

is sufficient to avoid that the expectation value of the χ^2 overestimates its true value by an amount larger than its statistical fluctuation. The advantage of Eq. (16) with respect to Eq. (8) is to provide a stability criterion that does not depend on the unknown matrix of uncertainties A , but only on the Frobenius norm of the matrix of fluctuations F . This dependence can be easily modelled as we explain in the next section.

3 Regularising the matrix of uncertainties

In this section we devise a procedure to regularise the matrix of uncertainties in such a way that the χ^2 becomes insensitive to inaccuracies in the estimation of the experimental uncertainties. We then demonstrate the effectiveness of the procedure in a toy model that is representative of realistic LHC measurements.

3.1 Regularisation procedure

Our aim is to obtain a regularised matrix of uncertainties A_{reg} which, for a given model of instabilities, fulfills the following criteria: *i*) A_{reg} is more stable than \bar{A} ; *ii*) A_{reg} is compatible with \bar{A} within the precision with which this is determined; and *iii*) the uncertainty estimated by A_{reg} never decreases in comparison to that estimated by \bar{A} . To this purpose, we first need to characterise the inaccuracies in the matrix \bar{A} , by means of a simplified model that builds upon the stability criterion, Eq. (16). We note that sometimes such a characterisation comes as part of the measurement itself, generally as the result of a dedicated analysis. In these cases, this characterisation has to be preferred to the model discussed below.

The model of inaccuracies that we devise ought to be minimal, general, and realistic. Minimal, because it should alter the matrix of uncertainties \bar{A} as little as possible; general, because it should be applied to any data set with no further information; and realistic, because it should capture the most likely sources of inaccuracy. These features lead us to making two assumptions.

The first assumption is that the correlations of experimental uncertainties across data points are determined much less precisely than the uncertainties for each data point, which we presume to be exact. This assumption is known to hold in practice, since the determination of certain correlations — such as those for two-point uncertainties defined as the difference between estimates obtained with two different Monte Carlo generators — require a certain amount of guesswork. This fact is occasionally reflected in different correlation models being presented with the measurement. Therefore we write

$$A = DA_{\text{corr}}, \tag{17}$$

where D is the $N_{\text{dat}} \times N_{\text{dat}}$ diagonal matrix of standard deviations for each data point

$$D_{ii} = \sqrt{\sum_j^{N_{\text{err}}} A_{ij}^2}. \tag{18}$$

We then assume that the covariance matrix provided by or built from the experiment has the true standard deviations, but correlations (encoded in \bar{A}_{corr} below) may be different

from the truth. Analogously to Eq. (9) we can therefore write

$$A = D(\bar{A}_{\text{corr}} + \delta F_{\text{corr}}). \tag{19}$$

Note that $A_{\text{corr}} A_{\text{corr}}^t$ is the covariance matrix of the reduced differences $(d_i - t_i)/D_{ii}$, hence the analysis carried out in Sect. 2 can be repeated *verbatim* for these variables. Analogously to Eq. (8), we can write

$$\Delta\chi^2 = \|A_{\text{corr}}^+ A_{\text{corr}}\| - N_{\text{dat}}, \tag{20}$$

and finally arrive at a stability criterion, similar to Eq. (16), under the assumption that D is well determined,

$$\|\bar{A}_{\text{corr}}^+\|_2 \|F_{\text{corr}}\|_F \leq \frac{1}{\sqrt{2}\delta}. \tag{21}$$

The second assumption is that $\|F_{\text{corr}}\|_F$ is independent of the number of data points or correlated experimental uncertainties in the measurement. The model then implies that the prevalent source of inaccuracy in the correlation matrix concentrates on a subset of data points and originates from a small number of correlated experimental uncertainties (for example the correlation of some two-point systematic uncertainties between the most extreme kinematic bins). While this assumption is a simplification, we find that the model is effective, as we will discuss in the context of both a toy model (see Sect. 3.2) and of a realistic case (see Sect. 4.2). If instead the source of inaccuracy in the correlation matrix arised from a number of systematic uncertainties that increased, e.g., with the number of data points N_{dat} , the regularisation procedure described below would over-regularise small data sets and under-regularise large ones when simultaneously applied to a collection of measurements.

Since F_{corr} is a matrix of adimensional coefficients (both units and magnitude of the data uncertainties are absorbed in D), we can simply set the norm to a constant, e.g. $\|F_{\text{corr}}\|_F = 1/\sqrt{2}$. Therefore, with the assumptions we have made, the model of uncertainties required to implement the stability criterion Eq. (7) contains one single adimensional parameter, δ , and the stability condition is

$$\|\bar{A}_{\text{corr}}^+\|_2 \leq \frac{1}{\delta}. \tag{22}$$

The free parameter δ characterises the precision of the correlation matrix. Its optimal value depends on the features of the measurement, and clearly cannot be obtained from the matrix itself. In the case of PDF determination, we will obtain it by studying the dependence of global fits on it, as we will discuss in Sect. 4.

The stability condition Eq. (22), together with the requirements presented at the beginning of this section, lay out a regularisation procedure. Specifically, Eq. (22) implies that the largest singular values of \bar{A}_{corr}^+ must be bound by δ^{-1} , and conversely that the smallest singular values of \bar{A}_{corr} must be bound by δ from below. The requirement that the regularised

matrix gives the same description as the original one in the directions that do not contribute to instability implies that the singular vectors with singular values greater than δ^{-1} are unchanged. Following Eq. (17), we write \bar{A} in terms of the singular value decomposition of \bar{A}_{corr} , $\bar{A}_{\text{corr}} = USV^t$

$$\bar{A} = DUSV^t, \tag{23}$$

and we can then define the regularised matrix A_{reg} as

$$\bar{A}_{\text{reg}} = DUS_{\text{reg}}V^t, \tag{24}$$

where S_{reg} is the matrix of singular values whose non-zero entries are

$$S_{\text{reg}(ii)} = \begin{cases} \delta & s_i < \delta \\ s_i & \text{otherwise} \end{cases}. \tag{25}$$

Note that, beside the formulation of the regularisation procedure laid out above, the stability condition, Eq. (22), can be inverted to quickly assess the stability of experimental uncertainties in a given measurement. We define the condition number Z as

$$Z = \|\bar{A}_{\text{corr}}^+\|_2 = \|\bar{A}_{\text{corr}}\|_2^{-1}. \tag{26}$$

It follows from Eqs. (22) and (25) that, if $Z > \delta^{-1}$, then it is likely that the precision with which correlations are determined is insufficient to ensure that they will not alter the expectation value of the χ^2 . This is demonstrated below with a toy model.

We note that the regularisation procedure can apply without modification to joint matrices constructed from multiple measurements when assuming the same value of δ for each of them. For example, if the measurements are independent, and the joint matrix is block diagonal, with each block corresponding to the covariance matrix from one measurement, the effect of the regularisation on the joint matrix is the same as applying it independently to each of the individual matrices, while the Z condition number will be the maximum across the measurements. Systematic uncertainties that are shared between measurements (hence making the joint matrix not completely block diagonal) also require no change in the procedure.

Finally, we remark that Eq. (25) should work for any value of δ , even if it has been derived by neglecting terms of $\mathcal{O}(\delta^2)$ in Eq. (12). This neglect, however, may make the interpretation of δ^{-1} as a measure of the precision with which correlations need to be known to ensure the stability of the χ^2 looser for small values of δ .

3.2 Toy model

We now apply the regularisation procedure devised in Sect. 3.1 to a toy model which is representative of a realistic LHC data set. This exercise will show how inaccuracies

in the degree of correlation of uncertainties can undermine the reliability of the χ^2 as a figure of merit.

The model consists of a data set made of four experimental data points, with a small uncorrelated statistical uncertainty of size ϵ , equal for each data point, and one correlated systematic uncertainty of size 1, affecting only the first three data points. The fourth point also has a systematic uncertainty, whose correlation with the other points is, however, not precisely known. We parametrise this lack of knowledge in terms of the variable x , and write the systematic uncertainty on the fourth point as a fluctuation, by an amount x , with respect to the other systematic uncertainty of size 1. We assume that the total variance is known. Note that this is consistent with the assumptions made in Sect. 3.1: correlations can fluctuate (in a way that, in the model, is parametrised by x), while variances remain fixed.

The matrix of uncertainties describing this toy model is

$$A(x) = \begin{pmatrix} \epsilon & 0 & 0 & 0 & 1 & 0 \\ 0 & \epsilon & 0 & 0 & 1 & 0 \\ 0 & 0 & \epsilon & 0 & 1 & 0 \\ 0 & 0 & 0 & \epsilon & 1-x & \sqrt{1-(1-x)^2} \end{pmatrix}. \tag{27}$$

By fixing the variance due to the systematics to 1 we let the parameter $\epsilon \ll 1$ control the relative size of the uncorrelated to correlated uncertainties. The parameter x can take values in the interval $[0, 2]$: $x = 0$ corresponds to the case in which the systematic uncertainty on the fourth data point is fully correlated with that of the other data points; $x = 1$ corresponds to the case of full decorrelation; and $x = 2$ corresponds to the case of full anti-correlation.

We now consider the situation in which the correlation is (inaccurately) estimated to be maximal, that is $x = \bar{x} = 0$. This inaccuracy is encoded in the matrix of uncertainties $\bar{A} = A(\bar{x})$:

$$\bar{A} = \begin{pmatrix} \epsilon & 0 & 0 & 0 & 1 & 0 \\ 0 & \epsilon & 0 & 0 & 1 & 0 \\ 0 & 0 & \epsilon & 0 & 1 & 0 \\ 0 & 0 & 0 & \epsilon & 1 & 0 \end{pmatrix}. \tag{28}$$

According to Eq. (6), the expectation value of the χ^2 given \bar{A} is

$$\langle \bar{\chi}^2 \rangle(x) = \|\bar{A}^+ A(x)\|_F^2 = 4 + \frac{6x}{\epsilon^2(\epsilon^2 + 4)}, \tag{29}$$

which has to be compared with the true expectation value given $A(x)$, see Eq. (5):

$$\langle \chi_{\text{true}}^2 \rangle(x) = \|A^+ A(x)\|_F^2 = N_{\text{dat}} = 4. \tag{30}$$

The situation is depicted in Fig. 1, where the curves obtained with either Eq. (29) or Eq. (30) are contrasted as a function of the true (unknown) variable x . We consider two illustrative values of the model parameter ϵ , 0.1 and 0.25, that correspond to the situation in which the uncorrelated

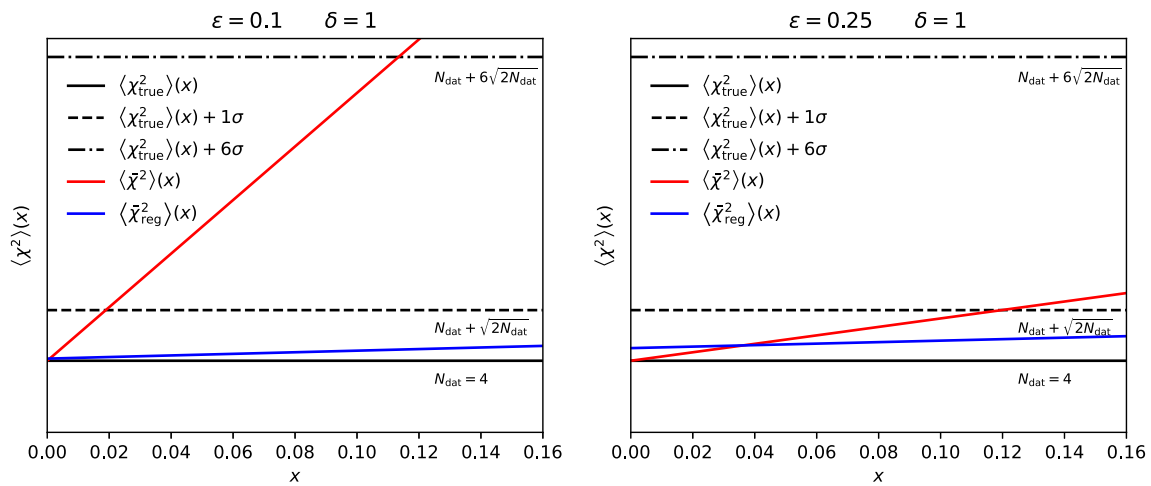


Fig. 1 The expectation value of the χ^2 , as a function of the variable x , in the toy model, for two values of the parameter ϵ : $\epsilon = 0.1$ (left) and $\epsilon = 0.25$ (right). We show: the true expectation value of Eq. (30), $\langle \chi_{\text{true}}^2 \rangle(x)$ (with one and six standard deviations for reference), given the matrix of uncertainties A , Eq. (27); the expectation value of Eq. (29),

$\langle \bar{\chi}^2 \rangle$, given the inaccurate matrix of uncertainties \bar{A} , Eq. (28); and the expectation value of Eq. (36), $\langle \bar{\chi}_{\text{reg}}^2 \rangle$, given the matrix of uncertainties \bar{A}_{reg} , Eq. (34), obtained after applying the regularisation procedure with $\delta = 1$

statistical uncertainty is equal, respectively, to 10% or 25% of the correlated systematic uncertainty. These values reflect the relative ratio of uncorrelated to correlated uncertainties in realistic current and future LHC measurements. As is apparent from Fig. 1, the incorrect estimation of \bar{x} leads to a large deviation of the expectation value of the χ^2 from its true value. The smaller the value of ϵ , the larger the deviation. For example, for a value of ϵ equal to 0.25, it is sufficient that the true value of x is 0.12 instead of zero to run afoul of the stability criterion of Eq. (7). For $\epsilon = 0.1$, the true value of x can be as small as 0.02 to encounter a similar instability.

We now apply the regularisation procedure devised in Sect. 3.1. We first write the matrix \bar{A} , Eq. (28), in terms of the matrices D and \bar{A}_{corr} , as per Eq. (17), which read

$$D = \sqrt{1 + \epsilon^2} I_{4 \times 4} \quad \text{and} \quad \bar{A}_{\text{corr}} = \frac{1}{\sqrt{1 + \epsilon^2}} \bar{A}. \quad (31)$$

The matrix of singular values for \bar{A}_{corr} is

$$S = \frac{1}{\sqrt{1 + \epsilon^2}} \begin{pmatrix} \epsilon & 0 & 0 & 0 & 0 & 0 \\ 0 & \epsilon & 0 & 0 & 0 & 0 \\ 0 & 0 & \epsilon & 0 & 0 & 0 \\ 0 & 0 & 0 & \sqrt{4 + \epsilon^2} & 0 & 0 \end{pmatrix}. \quad (32)$$

We denote the first three singular values as $s_{1,2,3} = \epsilon/\sqrt{1 + \epsilon^2}$ and the fourth one as $s_4 = \sqrt{4 + \epsilon^2}/\sqrt{1 + \epsilon^2}$, and note that $0 < s_{1,2,3} < s_4 < 2$ for any value of $\epsilon > 0$. We then apply the regularisation prescription given by Eqs. (23)–(25), by choosing $s_{1,2,3} < \delta^{-1} < s_4$. The regularised matrix of singular values therefore reads

$$S_{\text{reg}} = \begin{pmatrix} \delta^{-1} & 0 & 0 & 0 & 0 & 0 \\ 0 & \delta^{-1} & 0 & 0 & 0 & 0 \\ 0 & 0 & \delta^{-1} & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{\sqrt{4 + \epsilon^2}}{\sqrt{1 + \epsilon^2}} & 0 & 0 \end{pmatrix}, \quad (33)$$

and the regularised matrix of uncertainties

$$\bar{A}_{\text{reg}} = \begin{pmatrix} a & b & b & b & 1 & 0 \\ b & a & b & b & 1 & 0 \\ b & b & a & b & 1 & 0 \\ b & b & b & a & 1 & 0 \end{pmatrix}, \quad (34)$$

where

$$a = \frac{1}{4} (\epsilon + 3\delta^{-1} \sqrt{1 + \epsilon^2}) \quad \text{and} \quad b = \frac{1}{4} (\epsilon - \delta^{-1} \sqrt{1 + \epsilon^2}). \quad (35)$$

The expected value of the χ^2 is finally

$$\langle \bar{\chi}_{\text{reg}}^2 \rangle(x) = \left\| \bar{A}_{\text{reg}}^+ A(x) \right\|_F^2 = 4 + 6x \left(\frac{\delta^2}{1 + \epsilon^2} - \frac{1}{4 + \epsilon^2} \right) + \frac{12 \delta^2 \epsilon^2}{1 + \epsilon^2}. \quad (36)$$

The expression in Eq. (36) is compared to those in Eqs. (29)–(30) in Fig. 1 for the value $\delta = 1$. We note that this value fulfils the requirement $s_{1,2,3} < \delta^{-1} < s_4$ for any value of the parameter ϵ . As is apparent from Fig. 1, the regularisation procedure successfully achieves the goal for which it was devised: the expectation value of the regularised χ^2 , $\langle \chi_{\text{reg}}^2 \rangle$, does not differ from the true expectation value, $\langle \chi^2 \rangle_{\text{true}}$, by more than one standard deviation of the χ^2 distribution for any value of x .




Fig. 2 Stability of the toy model with additional assumptions on the value of the correlation. The green curve shows the deviation in χ^2 averaged over the assumed prior of the x parameter, Eq. (39). The black dashed horizontal line marks the limit from the stability criterion Eq. (7). The orange dashed vertical line, in the intersection, shows the most likely value of x that fulfills the stability criterion, Eq. (40)

The optimal value of δ should be determined on a case-by-case basis depending on the precision with which x is known. This is the topic that we will investigate in the next section in the context of PDF determination.

We now turn our attention to the situation where we can make further assumptions on the uncertainties in the determination of the correlation structure, for example when having access to additional information during the experimental analysis. In that case it might be advisable to study the effects on stability of various modelling choices, and the corresponding regularisation, in a more refined way than the one described in Sect. 3.1, where we strived for generality. We simulate this situation by assuming a specific prior for the value of the x parameter. We choose that prior to be a beta distribution with support in $x \in [0, 1]$ and such that $x = 0$ is the mode value. Specifically,

$$x \sim \text{Beta}(1, 5), \tag{37}$$

which corresponds to the probability density

$$f_x(\xi) = 5(1 - \xi)^4. \tag{38}$$

Our discussion implies that even though $x = 0$ is the most likely value, analyses using it are subject to instabilities. We can quantify this by computing the expected error in the χ^2 we would incur when assuming a particular value of x and when averaging over the distribution of possible values:

$$\langle \Delta\chi^2 \rangle(x) = \int_0^1 \left| \|\bar{A}^+(\xi)A(x)\|_F^2 - N \right| f_x(\xi) d\xi. \tag{39}$$

We represent $\langle \Delta\chi^2 \rangle_x(x)$ in Fig. 2, where we have set $\epsilon = 0.1$. The comparison with the limit imposed by the stability criterion Eq. (7), also displayed in Fig. 2, shows that presenting the covariance matrix with values too close

to the most likely value of the correlation under the prior yields large instabilities that would hamper the subsequent analysis. Selecting the most likely value that satisfies the stability criterion

$$x^* = \underset{\xi: \langle \Delta\chi^2 \rangle(\xi) \leq \sqrt{2N}}{\text{argmax}} f_x(\xi) \tag{40}$$

may be a way to decide the value of the correlation with which to present the covariance matrix. This would correspond to $x \approx 0.04$ under the settings presented here. Note that this small correction is consistent with the assumed knowledge of x , Eq. (37), but it would notably increase the accuracy of χ^2 computations using the covariance matrix.

The obvious disadvantage of this analysis is the difficulty of obtaining estimates for the covariance matrix parameters such as Eq. (37). These are unattainable outside the experimental collaborations responsible for the analysis and presumably challenging within. However, it may be useful to assess and refine correlation models internally. The regularisation procedure presented in Sect. 3 and applied to the toy model in Eq. (36) is indicated for the more common situation where such detailed information is missing. We demonstrate its usage for the problem of PDF determination next.

4 Determining PDFs with a regularised data set

In this section, we apply the regularisation procedure devised in Sect. 3 to a data set utilised for PDF determination. This is a particular problem relevant to LHC precision physics that relies on the χ^2 as a figure of merit. We first discuss how the regularisation procedure can be applied to characterise the data set that enters a given PDF determination. We then show how PDFs change if the nominal data set is replaced by a suitably regularised one, and study their dependence on the regularisation parameter δ . We finally investigate how the regularisation procedure performs in comparison to the correlation models provided with the measurements in the few cases in which these are available. All of our investigations are performed in the framework of the recent NNPDF4.0 PDF determination [9].

4.1 Characterising the data set

The NNPDF4.0 data set is the widest data set used for PDF determination to date. It consists of legacy fixed-target and collider deep-inelastic scattering and fixed-target Drell–Yan measurements, and of a wide range of measurements for various production processes in proton–proton collisions at the LHC. These include both Run I and Run II measurements and make about 30% of the NNPDF4.0 data set. Experimental uncertainties are typically of the order of few percent, the largest part of which is made of correlated systematic uncer-

tainties. A detailed description of the NNPDF4.0 data set is provided in Sect. 2 of [9].

Here we take a closer look at the LHC measurements that are part of the NNPDF4.0 data set, and in particular scrutinise the matrix of uncertainties of each measurement that contains more than one data point. The goal is to identify the measurements for which an inaccurate estimation of experimental correlations may significantly affect their χ^2 . To this purpose, for each measurement, we compute the condition number Z , Eq. (26), apply the regularisation procedure delineated in Sect. 3 for different values of the parameter δ , and evaluate how much the regularised covariance matrix differs from the nominal one. This piece of information is collected in Table 1, where we indicate, for each LHC measurement included in the NNPDF4.0 data set, its reference and the condition number Z ; we also indicate the maximum relative difference of the variances $\Delta\sigma_r$ and the maximum absolute difference of the correlation $|\Delta\rho|$ computed between the nominal data set and the data set regularised with $\delta^{-1} = 1, 2, 3, 4, 5, 7$. Blank spaces indicate that $\Delta\sigma_r = |\Delta\rho| = 0$, that is the regularisation procedure does not alter the nominal covariance matrix. We make two remarks.

First, one can single out the data sets for which an inaccurate estimation of experimental correlations may be of concern in a PDF fit. These are the data sets with the largest value of the condition number Z . If these data sets turn out to also have an unsatisfactory χ^2 in the fit, then additional investigations are needed to establish whether this is due solely to inaccurate experimental correlations, solely to inaccurate theoretical predictions, or to a combination of both. Conversely, if a data set has a low condition number Z but a large value of the χ^2 , the large value of the χ^2 is reasonably due to genuine inconsistencies between the data set and theory predictions. These considerations may help determine the optimal data set utilised as input to PDF determination, as done for the NNPDF4.0 parton set (see in particular Sect. 4.2 in [9]).

Second, one can determine the optimal value of the regularisation parameter δ in such a way that variances and correlations are not modified too much by the regularisation procedure in comparison to the nominal values. In this respect, inspection of Table 1 reveals that regularising the NNPDF4.0 data set with $\delta^{-1} = 1$ or $\delta^{-1} = 2$ is too aggressive, in that it leads to an increase of variances by an amount between 10% and 90%, and a variation of correlations between 0.1 and 0.5, depending on the data set. These figures are reduced, respectively, below 10% and 0.1 for $\delta^{-1} = 3$ and even further, to a few percent and below 0.05 for $\delta^{-1} = 4$ and $\delta^{-1} = 5$. Higher values of δ alter the nominal data set only minimally. As expected, the data sets associated to the highest condition number Z are those that are generally most affected by the regularisation procedure, in that they display the largest variation of variances and correlation; they also remain sensitive

to the regularisation procedure even if a modest amount of regularisation (that is, a high value of δ^{-1}) is applied. In the next section we shall see how these variations affect a fit of PDFs.

Among all of the LHC data sets collected in Table 1, we single out the two measurements that are associated to large values of Z and χ^2 (see Table 2) at the same time: ATLAS W, Z 7 TeV CC [15] and ATLAS dijets R=0.6 7 TeV [24]. They are representative of extreme cases in which small inaccuracies in the determination of experimental correlations can have a large impact on the computation of the χ^2 . Indeed these data sets have been the subject of much scrutiny [9, 15, 44, 45]. With a value of Z of order 10, it means that correlations must be estimated with an absolute uncertainty of roughly less than 0.1 in order to ensure that they make the χ^2 fluctuate by less than one standard deviation. If the correlation between two bins is estimated to be 1.0 while its real value is instead 0.9, one can expect the χ^2 to deviate significantly (by more than one standard deviation) from unity, even if there is good consistency between experimental central values and theoretical expectations.

Note that other data sets may have a large value of Z , e.g. CMS Z p_T 8 TeV [32], but not an anomalously large χ^2 (see Table 2). While our decorrelation procedure will also affect these data sets, as seen in Table 1, we do not consider them in the following discussion.

In Fig. 3 we show how the regularisation procedure described in Sect. 3.1 affects the covariance and correlation matrices of the two data sets singled out above. Specifically, we show the relative difference of the covariance matrix $\Delta\sigma_r$ and the difference of the correlation matrix $\Delta\rho$ for each of their elements, computed between the nominal data sets and the data set regularised with $\delta^{-1} = 4$. For ATLAS W, Z 7 TeV CC, we indicate the bins, differential in the rapidity of the lepton, η , corresponding to W^+ , W^- and Z production (the latter in three kinematic regions); for ATLAS dijets R=0.6 7 TeV, we indicate the bins, differential in the invariant mass of the dijet, m_{12} , corresponding to the six measured intervals of the absolute rapidity difference of the two leading jets, $|y^*|$: $0.0 \leq |y^*| \leq 0.5$; $0.5 \leq |y^*| \leq 1.0$; $1.0 \leq |y^*| \leq 1.5$; $1.5 \leq |y^*| \leq 2.0$; $2.0 \leq |y^*| \leq 2.5$; and $2.5 \leq |y^*| \leq 3.0$. As already noted, differences are small and do not exceed 5% for variances and 0.05 for correlations. These variations seem very reasonable to us; their effect, as well as that induced by larger (smaller) variations corresponding to more (less) aggressive regularisation will be investigated next.

4.2 Fitting PDFs

We now study the sensitivity of PDF determination to the regularisation procedure. To this purpose, we perform a series of fits, all based on the experimental, theoretical, and method-

Table 1 The LHC measurements included in the NNPDF4.0 data set [9]. For each measurement we indicate its reference, the condition number Z of the corresponding experimental covariance matrix, Eq (26), and the maximum relative difference of the variances $\Delta\sigma_r$ (in percent) and the maximum absolute difference of the correlation $|\Delta\rho|$ computed between the nominal data set and the data set obtained by applying the

regularisation procedure delineated in Sect. 3 for $\delta^{-1} = 1, 2, 3, 4, 5, 7$. Blank spaces indicate that $\Delta\sigma_r = |\Delta\rho| = 0$, that is the regularisation procedure does not alter the nominal covariance matrix. For ATLAS W, Z 7 TeV, CC and CF stand, respectively, for central-central and central-forward rapidity selections. We omit the data sets with a single data point

Data set	Ref.	Z	$\delta^{-1} = 1$		$\delta^{-1} = 2$		$\delta^{-1} = 3$		$\delta^{-1} = 4$		$\delta^{-1} = 5$		$\delta^{-1} = 7$	
			$\Delta\sigma_r$	$ \Delta\rho $	$\Delta\sigma_r$	$ \Delta\rho $	$\Delta\sigma_r$	$ \Delta\rho $	$\Delta\sigma_r$	$ \Delta\rho $	$\Delta\sigma_r$	$ \Delta\rho $	$\Delta\sigma_r$	$ \Delta\rho $
ATLAS LM DY 7 TeV	[12]	3.7	71.0	0.51	12.0	0.15	2.51	0.03						
ATLAS HM DY 7 TeV	[13]	3.1	67.3	0.50	8.99	0.13	0.43	0.01						
ATLAS W, Z 7 TeV	[14]	3.5	87.2	0.45	15.2	0.13	2.36	0.02						
ATLAS W, Z 7 TeV CC	[15]	9.0	94.4	0.50	21.9	0.19	8.63	0.09	4.15	0.05	2.12	0.02	0.50	0.01
ATLAS W, Z 7 TeV CF	[15]	2.8	69.7	0.49	9.13	0.10								
ATLAS LM DY 2D 8 TeV	[16]	1.1	15.1	0.09										
ATLAS HM DY 2D 8 TeV	[17]	2.6	79.6	0.40	8.42	0.08								
ATLAS $\sigma_{W,Z}^{\text{tot}}$ 13 TeV	[18]	5.0	57.6	0.60	11.1	0.18	3.72	0.07	1.15	0.02				
ATLAS W^+ +jet 8 TeV	[19]	4.0	78.7	0.52	12.2	0.16	2.55	0.05	0.03					
ATLAS W^- +jet 8 TeV	[19]	5.7	82.0	0.54	14.9	0.18	4.74	0.07	1.78	0.03	0.51	0.01		
ATLAS Z $p_T, m_{\ell\ell}$ 8 TeV	[20]	3.3	71.6	0.44	7.99	0.09								
ATLAS Z p_T, y_Z 8 TeV	[20]	8.8	83.1	0.43	11.6	0.11								
ATLAS $t\bar{t}$ ℓ +jets y_t 8 TeV	[21]	1.6	37.3	0.34										
ATLAS $t\bar{t}$ ℓ +jets $y_{t\bar{t}}$ 8 TeV	[21]	2.2	52.5	0.41	2.31	0.03								
ATLAS $t\bar{t}$ 2ℓ $y_{t\bar{t}}$ 8 TeV	[22]	1.9	31.5	0.28										
ATLAS jets $R=0.6$ 8 TeV	[23]	5.5	93.6	0.48	20.6	0.17	7.21	0.07	2.61	0.03	0.53	0.01		
ATLAS dijets $R=0.6$ 7 TeV	[24]	10	95.7	0.51	22.3	0.19	9.08	0.09	4.54	0.05	2.53	0.03	0.80	0.01
ATLAS γ 13 TeV	[25]	1.3	38.2	0.14										
ATLAS single t dy_t 7 TeV	[26]	1.3	20.8	0.22										
ATLAS single t y_t 7 TeV	[26]	1.4	29.7	0.25										
ATLAS single t dy_t 8 TeV	[27]	1.2	13.5	0.12										
ATLAS single t y_t 8 TeV	[27]	1.2	15.6	0.17										
CMS W e asy. 7 TeV	[28]	1.0	9.17	0.03										
CMS W μ asy. 7 TeV	[29]	1.2	30.3	0.12										
CMS DY 2D 7 TeV	[30]	8.8	85.2	0.52	17.0	0.18	5.99	0.08	2.86	0.04	1.53	0.02	0.42	0.01
CMS W rapidity 8 TeV	[31]	13	93.9	0.59	22.6	0.21	9.39	0.10	4.88	0.05	2.79	0.03	1.03	0.01
CMS Z p_T 8 TeV	[32]	9.5	87.7	0.46	15.9	0.14	3.09	0.03						
CMS dijets 7 TeV	[33]	4.7	88.8	0.48	18.2	0.15	5.51	0.05	1.22	0.01				
CMS jets 8 TeV	[34]	6.3	92.4	0.53	20.2	0.19	7.14	0.08	2.78	0.03	0.93	0.01		
CMS $t\bar{t}$ ℓ +jets 8 TeV	[35]	1.6	41.9	0.26										
CMS $t\bar{t}$ 2D 2ℓ 8 TeV	[36]	1.9	58.3	0.33	5.84	0.06								
CMS $t\bar{t}$ 2ℓ 13 TeV	[37]	5.2	76.6	0.48	11.5	0.13	2.11	0.03	0.51	0.01	0.05			
CMS $t\bar{t}$ ℓ +jet 13 TeV	[38]	7.5	83.3	0.51	17.7	0.17	6.36	0.07	2.67	0.03	1.06	0.01	0.07	
LHCb $Z \rightarrow ee$ 7 TeV	[39]	1.4	55.1	0.34										
LHCb $W, Z \rightarrow \mu$ 7 TeV	[40]	2.9	66.6	0.40	5.22	0.08								
LHCb $Z \rightarrow ee$ 8 TeV	[41]	1.4	45.3	0.20										
LHCb $W, Z \rightarrow \mu$ 8 TeV	[42]	2.5	69.8	0.43	5.28	0.07								
LHCb $Z \rightarrow ee$ 13 TeV	[43]	2.4	54.9	0.26										
LHCb $Z \rightarrow \mu\mu$ 13 TeV	[43]	1.6	74.9	0.39	5.69	0.06								

Table 2 The number of data points, N_{dat} , and the χ^2 per data point, χ^2/N_{dat} , for the NNPDF4.0 NNLO baseline fit and for each of the fits performed with the regularisation procedure delineated in Sect. 3 for $\delta^{-1} = 1, 2, 3, 4, 5, 7$

Data set	N_{dat}	χ^2/N_{dat}						
		NNPDF4.0	$\delta^{-1} = 1$	$\delta^{-1} = 2$	$\delta^{-1} = 3$	$\delta^{-1} = 4$	$\delta^{-1} = 5$	$\delta^{-1} = 7$
Deep-inelastic scattering	3089	1.12	0.64	1.02	1.09	1.11	1.12	1.12
Fixed-target Drell–Yan	195	0.98	0.48	0.90	0.96	0.97	0.97	0.99
Tevatron Drell–Yan	65	1.11	0.48	0.71	0.85	0.93	1.02	1.10
ATLAS total	679	1.24	0.50	0.84	0.97	1.04	1.10	1.19
LM DY 7 TeV	6	0.88	0.22	0.60	0.85	0.88	0.88	0.88
HM DY 7 TeV	5	1.69	0.44	1.25	1.64	1.69	1.69	1.69
W, Z 7 TeV	30	0.98	0.24	0.65	0.96	1.01	1.01	1.00
W, Z 7 TeV CC	46	1.92	0.31	0.74	0.94	1.21	1.47	1.76
W, Z 7 TeV FC	15	1.03	0.48	0.96	1.04	1.04	1.04	1.04
LM DY 2D 8 TeV	48	1.11	0.61	1.09	1.12	1.11	1.11	1.12
HM DY 2D 8 TeV	60	1.22	1.18	1.21	1.22	1.22	1.22	1.22
$\sigma_{W,Z}^{\text{tot}}$ 13 TeV	3	0.77	0.30	0.39	0.46	0.62	0.79	0.77
W^+ +jet 8 TeV	15	0.79	0.45	0.58	0.70	0.79	0.79	0.79
W^- +jet 8 TeV	15	1.49	0.79	1.15	1.29	1.38	1.44	1.49
$Z p_T, m_{\ell\ell}$ 8 TeV	44	0.90	0.43	0.86	0.90	0.91	0.90	0.90
$Z p_T, y_Z$ 8 TeV	48	0.90	0.22	0.65	0.91	0.90	0.90	0.90
$\sigma_{t\bar{t}}$ 7, 8, 13 TeV	3	1.64	1.47	1.70	1.74	1.79	1.77	1.70
$t\bar{t}$ ℓ +jets y_t 8 TeV	4	3.28	1.42	2.45	2.92	2.98	3.01	3.09
$t\bar{t}$ ℓ +jets $y_{t\bar{t}}$ 8 TeV	4	3.83	1.00	2.62	3.32	3.40	3.46	3.52
$t\bar{t}$ $2\ell y_{t\bar{t}}$ 8 TeV	5	1.62	0.65	1.46	1.55	1.57	1.59	1.59
jets $R=0.6$ 8 TeV	171	0.68	0.38	0.60	0.68	0.69	0.69	0.69
dijets $R=0.6$ 7 TeV	90	2.14	0.23	0.59	0.85	1.10	1.38	1.87
γ 13 TeV	53	0.76	0.58	0.72	0.76	0.77	0.76	0.76
single $t R_t$ 7, 13 TeV	2	0.28	0.24	0.27	0.28	0.28	0.28	0.28
single $t dy_t$ 7 TeV	3	0.96	0.94	0.98	0.97	0.96	0.96	0.96
single $t y_t$ 7 TeV	3	0.06	0.03	0.06	0.06	0.06	0.06	0.06
single $t dy_t$ 8 TeV	3	0.25	0.20	0.23	0.24	0.24	0.24	0.24
single $t y_t$ 8 TeV	3	0.19	0.17	0.19	0.19	0.19	0.19	0.19
CMS total	474	1.31	0.39	0.83	1.08	1.21	1.26	1.28
$W e$ asy. 7 TeV	11	0.84	0.71	0.79	0.79	0.79	0.79	0.81
$W \mu$ asy. 7 TeV	11	1.70	1.34	1.76	1.76	1.76	1.76	1.75
DY 2D 7 TeV	110	1.36	0.43	1.05	1.27	1.33	1.35	1.36
W rapidity 8 TeV	22	1.33	0.13	0.16	0.22	0.33	0.46	0.80
$Z p_T$ 8 TeV	28	1.40	0.21	0.75	1.25	1.40	1.40	1.40
dijets 7 TeV	54	1.79	0.63	1.40	1.73	1.84	1.84	1.82
jets 8 TeV	185	1.19	0.24	0.53	0.84	1.07	1.15	1.18
$\sigma_{t\bar{t}}$ 5, 8, 7, 13 TeV	4	0.45	0.39	0.49	0.51	0.53	0.52	0.48
$t\bar{t}$ ℓ +jets 8 TeV	9	1.25	0.67	1.13	1.22	1.23	1.23	1.23
$t\bar{t}$ 2D 2 ℓ 8 TeV	16	1.01	0.41	0.86	1.07	1.07	1.06	1.04
$t\bar{t}$ 2 ℓ 13 TeV	10	0.52	0.09	0.26	0.37	0.43	0.50	0.51
$t\bar{t}$ ℓ +jet 13 TeV	11	0.64	0.04	0.15	0.33	0.42	0.51	0.59
single $t R_t$ 7, 8, 13 TeV	3	0.42	0.38	0.41	0.42	0.42	0.42	0.42
LHCb total	116	1.55	0.73	1.41	1.53	1.56	1.55	1.55

Table 2 continued

Data set	N_{dat}	χ^2/N_{dat}						
		NNPDF4.0	$\delta^{-1} = 1$	$\delta^{-1} = 2$	$\delta^{-1} = 3$	$\delta^{-1} = 4$	$\delta^{-1} = 5$	$\delta^{-1} = 7$
$Z \rightarrow ee$ 7 TeV	9	1.64	0.88	1.55	1.62	1.63	1.63	1.65
$W, Z \rightarrow \mu$ 7 TeV	29	1.94	0.66	1.61	1.93	1.97	1.97	1.96
$Z \rightarrow ee$ 8 TeV	17	1.34	0.88	1.30	1.31	1.33	1.33	1.33
$W, Z \rightarrow \mu$ 8 TeV	30	1.42	0.73	1.31	1.41	1.45	1.43	1.43
$Z \rightarrow ee$ 13 TeV	15	1.73	0.95	1.71	1.72	1.73	1.73	1.72
$Z \rightarrow \mu\mu$ 13 TeV	16	1.00	0.31	0.91	0.98	0.99	0.99	0.99
Total	4618	1.16	0.58	0.97	1.07	1.11	1.13	1.15

ological input that enters the default next-to-next-to-leading order (NNLO) NNPDF4.0 parton set (see [9] for details), in which we regularise the data set. Specifically, we perform six fits in each of which we consider a different amount of regularisation, namely $\delta^{-1} = 1, 2, 3, 4, 5, 7$. All the fits are made of $N_{\text{rep}} = 100$ Monte Carlo replicas. Note that these fits are different from those presented in Sect. 8.7 of [9]: here the regularisation procedure is applied to the NNPDF4.0 data set as a whole (and indeed to the total covariance matrix), while there it was applied only to a specific measurement (that was part of the NNPDF4.0 data set or not) at a time.

In Table 2 we display the value of the χ^2 per data point, χ^2/N_{dat} , for each of these fits, and compare it to that of the NNLO NNPDF4.0 default fit. Deep-inelastic scattering, fixed-target Drell–Yan, and Tevatron Drell–Yan measurements, which are mostly unaffected by the regularisation procedure, are all aggregated; ATLAS, CMS and LHCb measurements are instead displayed individually. The total values (for each experiment and for the total data set) are also shown, as well as the corresponding number of data points.

In Fig. 4 we then display the resulting PDFs, specifically the up, anti-up, down, anti-down, strange, anti-strange, charm and gluon PDFs at a scale $Q = 100$ GeV. PDFs are compared to the NNPDF4.0 NNLO baseline parton set, and are normalised to its central value. For $\delta^{-1} = 1, 3, 5, 7$ we display only the central value. Otherwise uncertainties correspond to 68% confidence levels.

A joint inspection of Table 2 and of Fig. 4 reveals some interesting features. We first observe that, as expected, the regularisation procedure has a significant effect on the χ^2 . A general decrease of its value is observed in comparison to NNPDF4.0, by an amount that increases with the increase in the amount of regularisation (that is, with the decrease of the value of δ^{-1}). For the largest value $\delta^{-1} = 7$, no statistically significant differences are seen with respect to NNPDF4.0, neither in the value of the χ^2 per data point nor in PDFs. Conversely, for small values of δ^{-1} , $\delta^{-1} = 1$ and $\delta^{-1} = 2$, the total χ^2 per data point drops from 1.16 to 0.58 and 0.97, respectively. These variations correspond to a 28σ and a 9σ

fluctuation in units of the χ^2 standard deviation, which obviously denote an excessive regularisation of the NNPDF4.0 data set. As noted at the end of Sect. 3, such an excessive regularisation may also arise from neglecting terms of $\mathcal{O}(\delta^2)$ in Eq. (12).

The PDFs obtained in the fit with $\delta^{-1} = 1$ (and similarly in the fit with $\delta^{-1} = 2$, which is not displayed in Fig. 4) are indeed consistently distorted in comparison to NNPDF4.0. The central value of the former fluctuates, in units of the NNPDF4.0 PDF uncertainty around the central value of the latter, by about one standard deviation for the up, anti-up, down and anti-down PDFs, and slightly more for the strange, anti-strange, charm and gluon PDF. In this respect, it is worth noting that the strange and gluon PDFs are sensitive, respectively, to the ATLAS W, Z 7 TeV CC [15] and ATLAS dijets $R=0.6$ 7 TeV [24] data sets: these have some of the largest values of Z and display the largest reduction of χ^2 upon regularisation.

The outer cases corresponding to $\delta^{-1} = 1, 2, 7$ are therefore to be interpreted as a validation of the regularisation procedure, which behaves as expected. The fits corresponding to $\delta^{-1} = 3, \delta^{-1} = 4$ and $\delta^{-1} = 5$ are instead more interesting. Variations of the χ^2 with respect to NNPDF4.0 correspond, respectively, to a 3.8σ , 2.4σ and 1.4σ fluctuation in units of the χ^2 standard deviation. Interestingly, the difference between the expected $\chi^2/N_{\text{dat}} = 1$ and the χ^2 obtained in the fits corresponding to $\delta^{-1} = 3, 4, 5$ amounts, respectively, to 3.3σ , 5.3σ and 6.2σ in units of the χ^2 standard deviation. This is a significant reduction in comparison to 7.7σ of the default NNPDF4.0 determination.

Such an improvement in the χ^2 statistic is accompanied by remarkably limited PDF variations if one compares the fits with $\delta^{-1} = 3, 4, 5$ with NNPDF4.0. Central values fluctuate by a small fraction of the NNPDF4.0 PDF uncertainty, except for the gluon PDF, which varies by up to half of the NNPDF4.0 uncertainty around $x \sim 0.3$; PDF uncertainties are almost unaffected. Remarkably, all these variations are much smaller than those due to variations of the data set itself (see Sect. 7 in [9]).

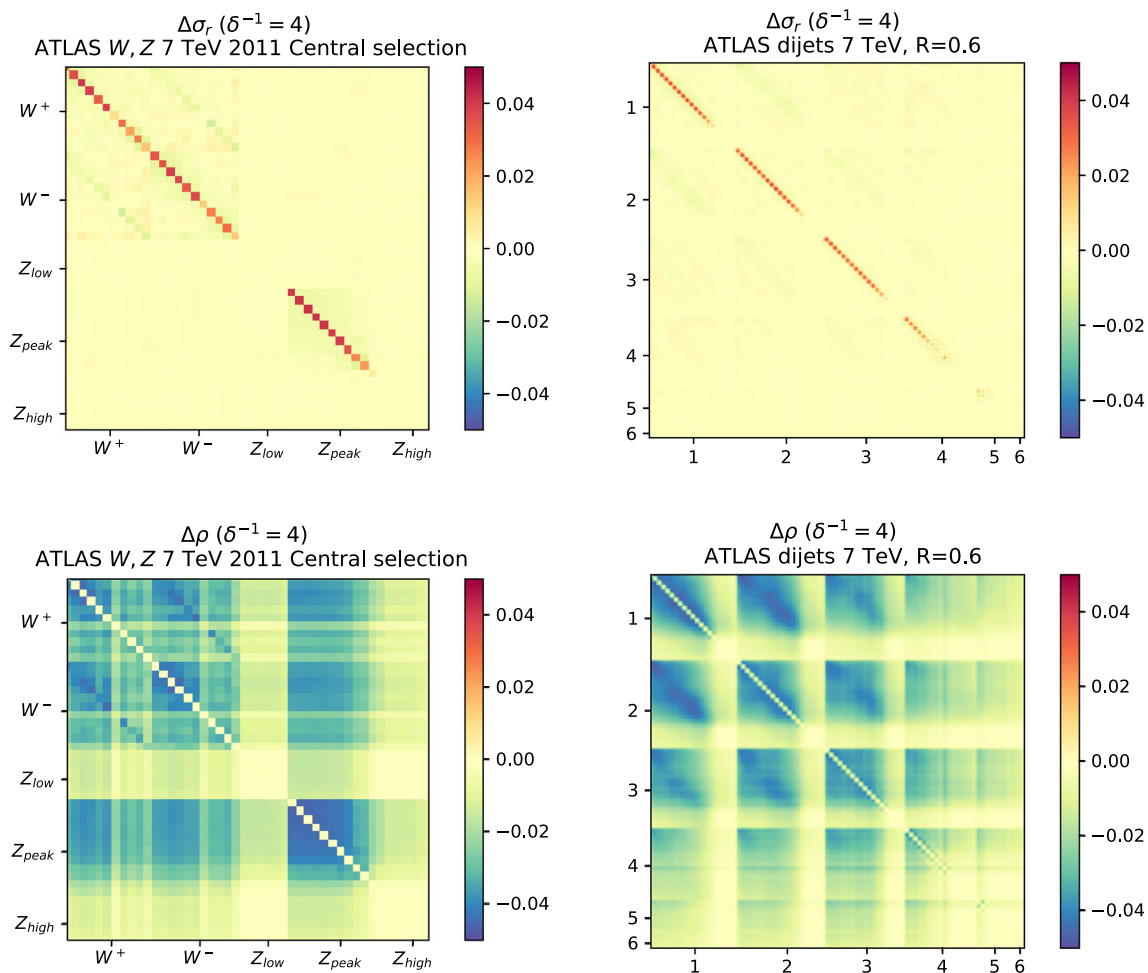


Fig. 3 The relative difference of the covariance matrix $\Delta\sigma_r$ (top) and the difference of the correlation matrix $\Delta\rho$ (bottom) for each of their elements, computed between the regularised data sets and the data set nominal with $\delta^{-1} = 4$. We show results for the two measurements in the NNPDF4.0 data set that have the largest value of Z , see Table 1: ATLAS W, Z 7 TeV CC [15] (left) and ATLAS dijets R=0.6 7 TeV [24] (right).

The fact that PDFs do not vary significantly in the fits to the regularised data set with $\delta^{-1} = 3, 4, 5$ is further displayed in Fig. 5, where we show a data–theory comparison for some selected bins of the ATLAS W, Z 7 TeV CC [15] and dijets R=0.6 7 TeV [24] measurements. Specifically, we show the W^+ and W^- subsets, as a function of the absolute value of the lepton rapidity η , for the former, and two bins in the absolute rapidity difference between the two leading jets $|y^*|$ as a function of the di-jet invariant mass m_{12} , for the latter. Theoretical predictions are obtained with the NNPDF4.0 baseline parton set and with the PDFs obtained by fitting the NNPDF4.0 data set regularised with $\delta^{-1} = 1, 3, 4, 5, 7$. They are all accurate to NNLO in the strong coupling, both in the PDFs and in the matrix elements. Results are shown as ratios to the experimental central value, with one-sigma experimental and PDF uncertainties.

For ATLAS W, Z 7 TeV CC, we indicate the bins, differential in the rapidity of the lepton, η , corresponding to W^+ , W^- and Z production (the latter in three kinematic regions); for ATLAS dijets R=0.6 7 TeV we indicate the bins, differential in the invariant mass of the dijet, m_{12} , corresponding to the six measured intervals of the absolute rapidity difference of the two leading jets, $|y^*|$, see text for details

The experimental uncertainty is the sum in quadrature of the statistical and of all systematic uncertainties.

As noted in Sect. 4.1, the data sets displayed in Fig. 5 are those with large values of Z and χ^2 , and for which the regularisation procedure introduces some of the largest differences in the variances and correlations of the data, see Table 1 and Fig. 3. In spite of this, only small differences are observed between predictions obtained with NNPDF4.0 and any of the regularised fits with $\delta^{-1} = 1, 3, 4, 5, 7$; slightly larger fluctuations are observed in the fit with a large amount of regularisation ($\delta^{-1} = 1$), albeit only for the data points at central rapidity, for ATLAS W, Z 7 TeV CC, or at large invariant mass, for ATLAS dijets R = 0.6 7 TeV.

We therefore conclude that the PDFs obtained from any of the regularised fits with $\delta^{-1} = 3, 4, 5$ represent the same underlying truth as the NNPDF4.0 parton set. They however

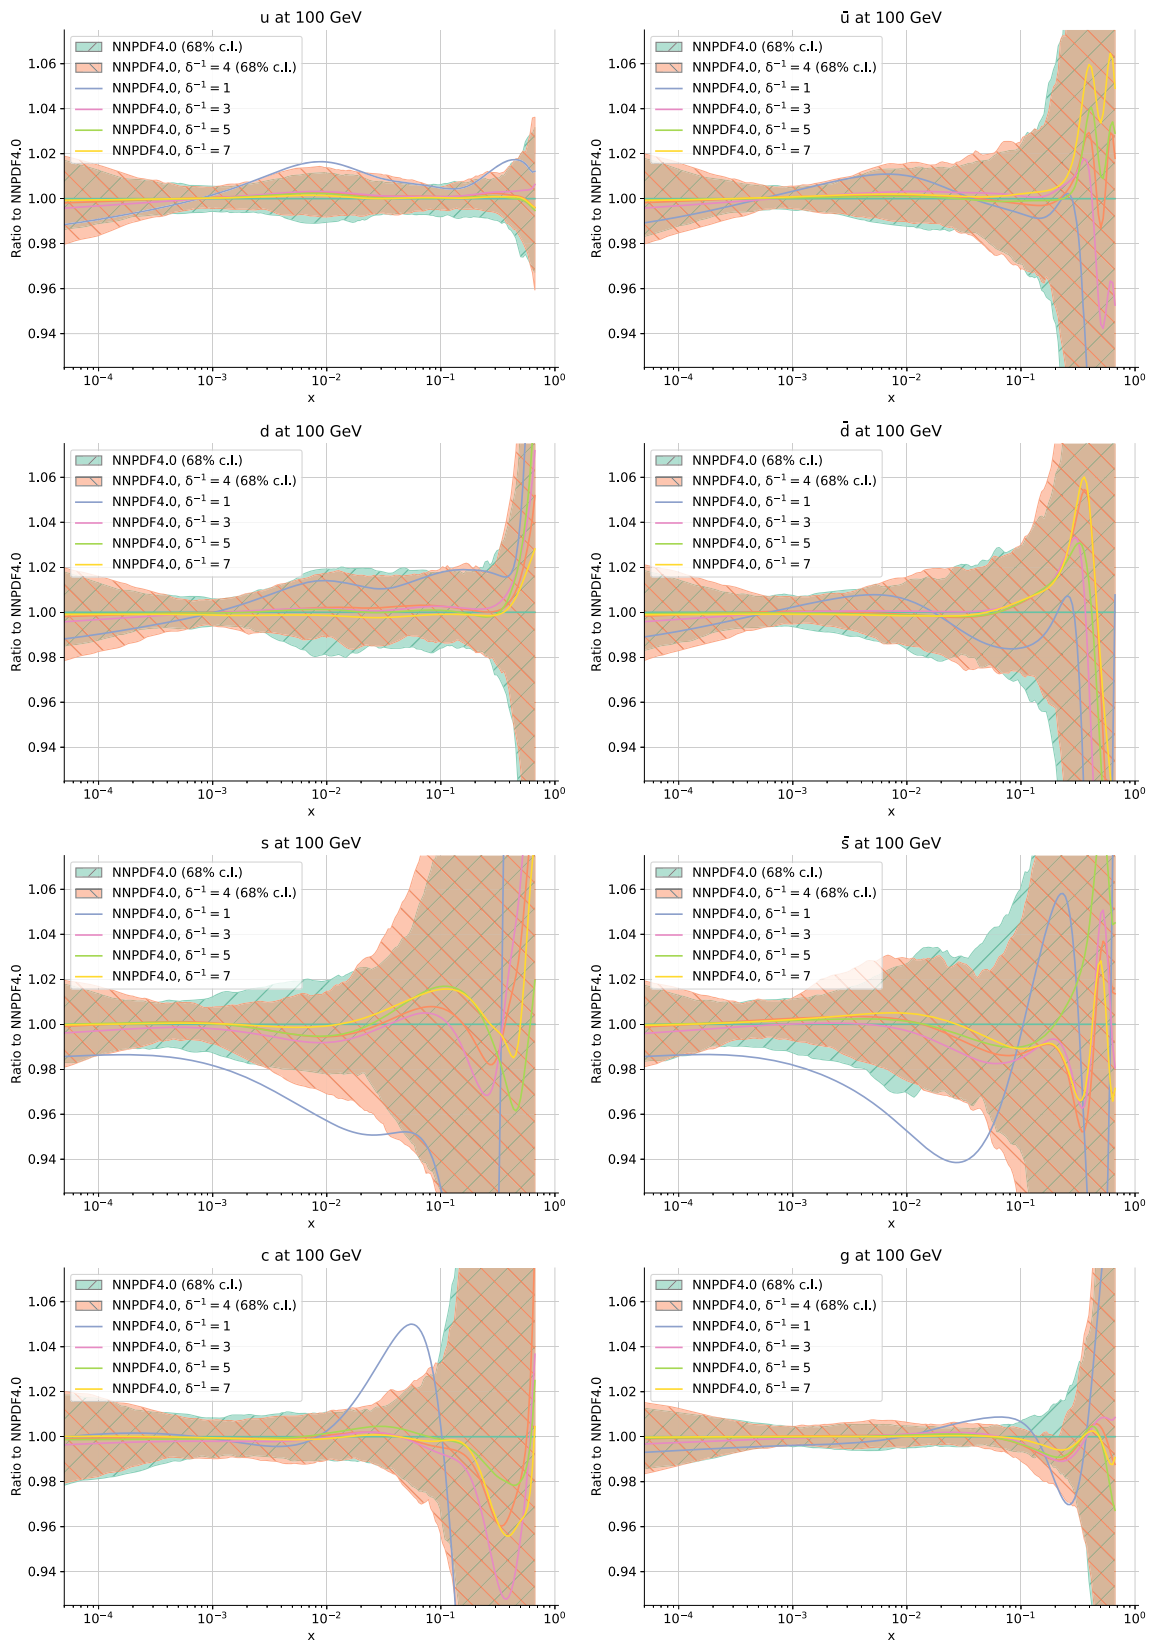


Fig. 4 The PDFs obtained by fitting the NNPDF4.0 data set after regularisation with different values of the parameter $\delta^{-1} = 1, 3, 4, 5, 7$. From top to bottom, left to right, we show the up, anti-up, down, anti-down, strange, anti-strange, charm and gluon PDFs at a scale

$Q = 100$ GeV. PDFs are compared to the NNPDF4.0 baseline parton set, and normalised to its central value. For $\delta^{-1} = 1, 3, 5, 7$ we display only the central value. Otherwise uncertainties correspond to 68% confidence levels. All PDF fits are accurate to NNLO

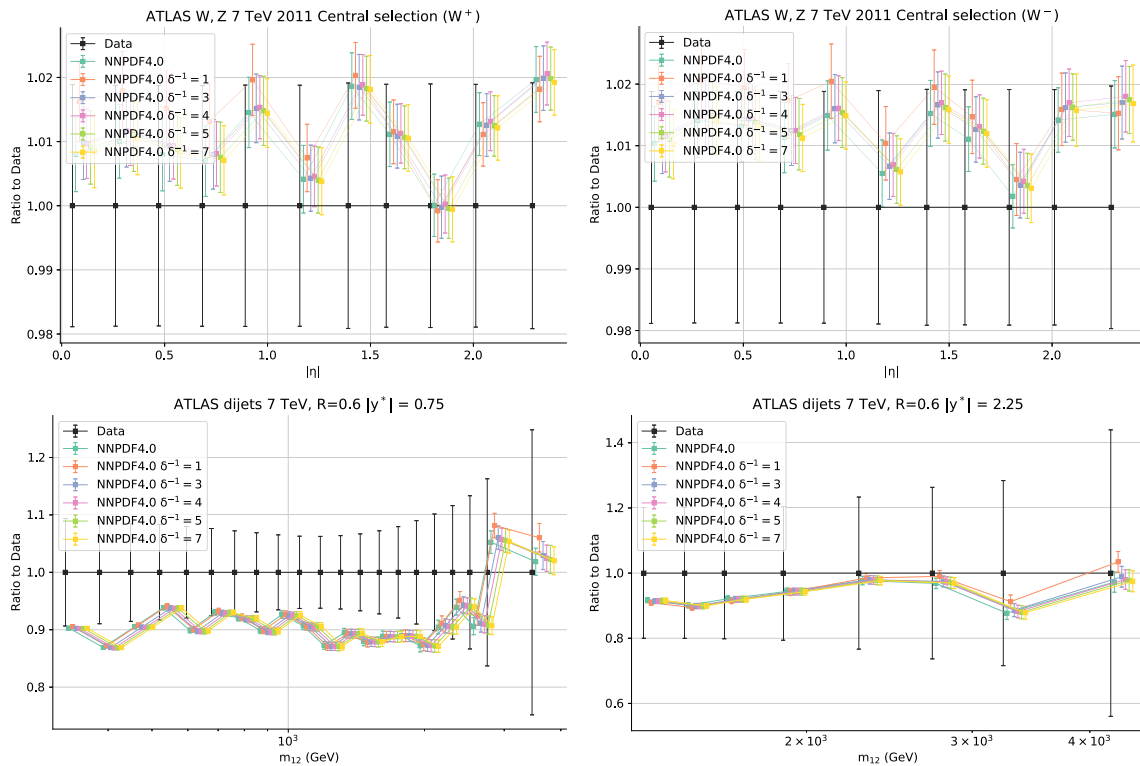


Fig. 5 Data–theory comparison for the W^\pm subset of the ATLAS W, Z 7 TeV CC measurement [15], as a function of the absolute lepton rapidity η (top), and for two bins in the absolute rapidity difference between the two leading jets $|y^*|$ of the ATLAS dijets $R=0.6$ 7 TeV measurement [24], as a function of the di-jet invariant mass m_{12} . Theoretical predictions are obtained with the NNPDF4.0 baseline parton

lead to a χ^2 that is better than the NNPDF4.0 one by up to 4σ , in units of the χ^2 standard deviation, and that is only about 3σ away from the expectation of unit χ^2 (instead of about 8σ). In other words, the nominal χ^2 determined in [9] is likely to be spuriously inflated by inaccuracies in the estimation of the experimental correlations in the LHC data. Further discrimination among the equally good values $\delta^{-1} = 3, 4, 5$ can be made on the basis of how big the changes to the covariance matrices are in relation to the precision at which they are estimated. Since the precision is unknown, this entails a degree of subjectivity. We deem that the values of $\Delta\sigma_r < 5\%$ and $|\Delta\rho| < 0.05$ implied by $\delta^{-1} = 4$ suggest it safe to assume that the resulting regularised covariance matrices are compatible with the original ones within the precision at which they were determined, while ensuring stability against possibly bigger inaccuracies in the correlations. Therefore, the fit with $\delta^{-1} = 4$ will be used as reference in the remainder of this paper.

set and with the PDFs obtained by fitting the NNPDF4.0 data set regularised with $\delta^{-1} = 1, 3, 4, 5, 7$. They are all accurate to NNLO in the strong coupling, both in the PDFs and in the matrix elements. Results are shown as ratios to the experimental central value, with one-sigma PDF and experimental uncertainties. The latter is the sum in quadrature of the statistical and of all systematic uncertainties

4.3 Correlating and decorrelating experimental uncertainties with more information

As we have mentioned in Sect. 3.1, the correlation models provided with the measurements have to be preferred to our regularisation procedure whenever these are available, and if they result in a stable covariance matrix. For example, the correlation model recommended in [23] for the analysis of the ATLAS jets $R=0.6$ 8 TeV measurement is used by default in the NNPDF4.0 determination [9] and in all the fits presented in Sect. 4.2. It is therefore not surprising that the regularisation procedure has almost no impact on the χ^2 of this specific data set.

Correlation models, which follow from a careful experimental analysis of all of the sources of systematic uncertainties and of their correlations, are however not always available. Sometimes they become available only long after the measurement is published, and sometimes a clear recommendation for their usage is not provided. In order to remedy this

lack of information, some guesswork is carried out to identify the systematic uncertainties whose nominal correlations are likely to be too strong. For instance, two of these [46, 47] have targeted, respectively, the ATLAS 7 TeV single-inclusive jet measurement [48] and the 8 TeV top-pair lepton+jet measurement [21]. They were performed in the framework of the MMHT2014 global analysis [49] by inspecting the nuisance parameters associated to each systematic uncertainty in the χ^2 . Similar studies [50, 51], targeting the same measurements and based on complete decorrelation of certain systematic uncertainties, were also carried out in the framework on the NNPDF3.1 global analysis [52]. Sometimes these analyses have been used to inform and/or validate the experimental correlation models. In this respect, our regularisation procedure can be utilised in the same spirit, with the advantage that it is more general and requires less information than the aforementioned analyses.

Here we investigate how the regularisation procedure performs in comparison to the correlation models provided with the measurement in the few cases in which these are available. We consider two cases. The first case concerns the ATLAS dijets R=0.6 7 TeV [24] measurement, for which a STRONG and a WEAK (de-)correlation models are provided on top of the nominal correlation model used in NNPDF4.0 and in all the fits of Sect. 4.2. None of these models are clearly recommended in [24], hence why they have not been previously considered. The second case concerns three ATLAS 8 TeV measurements, namely the W^\pm +jet [19], the $t\bar{t}$ ℓ +jets [21], and the single-inclusive jets R=0.6 [23] measurements. Details on how to correlate or decorrelate systematic uncertainties between bins within and across these measurements have been provided only very recently [53]. This is the reason why they have not been previously considered. We will refer to this correlation model with the label ATLAS henceforth.

We then perform four fits, all based on the experimental, theoretical, and methodological input that enters the default NNPDF4.0 parton set, by considering these correlation models. The first two fits are performed using, respectively, the STRONG and WEAK correlation models for the ATLAS dijets R = 0.6 7 TeV measurement. Experimental correlations for all of the other data sets are otherwise as in NNPDF4.0. The third fit is performed using the ATLAS correlation model for all the concerned ATLAS 8 TeV measurements. This correlation model was not completely utilised in NNPDF4.0 (in particular for what concerns correlations between pairs of points belonging to different data sets). It also does not enter the two aforementioned fits. The fourth fit is performed by combining the WEAK and ATLAS correlation models at the same time.

In Table 3 we display the value of the χ^2 per data point, χ^2/N_{dat} , for each of these fits and compare it to that of the NNLO NNPDF4.0 default fit, and of the fit obtained by regularising the NNPDF4.0 data set with $\delta^{-1} = 4$. For con-

ciseness, we aggregate the data sets into one of the following classes: deep-inelastic scattering, fixed-target Drell–Yan, Tevatron Drell–Yan, ATLAS, CMS, and LHCb. For ATLAS, we also indicate the individual χ^2 of the data sets affected by the correlation models. The corresponding number of data points, N_{dat} , is also indicated.

In Fig. 6 we show the resulting PDFs, specifically the anti-up, anti-down, charm and gluon PDFs at a scale $Q = 100$ GeV. PDFs are compared to the NNPDF4.0 NNLO baseline parton set, and to the PDFs obtained by regularising the NNPDF4.0 data set with $\delta^{-1} = 4$. All the curves are normalised to the NNPDF4.0 central value. For all PDFs but the NNPDF4.0 NNLO baseline, we show only the central value. Otherwise the uncertainty corresponds to the 68% confidence interval.

A joint inspection of Table 3 and Fig. 6 reveals two features. First, the fit quality, as quantified by the value of the χ^2 per data point, does not change upon variation of the available correlation models, either for the data sets affected by the model, or for the other data sets. This behaviour contrasts with the larger variations seen upon refitting a regularised data set, even when the amount of regularisation is fairly limited, see Table 2. Second, the shifts of PDF central values induced by a given correlation model, however modest they turn out to be, are generally very close to the shifts induced by the regularisation of the data set (specifically with $\delta^{-1} = 4$). This is apparent for the WEAK correlation model, whose only feature is to partially decorrelate certain uncertainties in one of the data sets with the largest condition number, ATLAS dijets R=0.6 7 TeV [24]. In this respect, this correlation model is relatively close to what the regularisation procedure achieves. For other correlation models, qualitatively similar shifts are also seen, although with some quantitative differences.

The fact that our regularisation procedure captures the same qualitative shifts on PDF central values as experimental correlation models, once one or the others are used to determine the PDFs, is suggestive. Whether this is a coincidental feature, limited to the correlation models considered, or a more general one, could only be investigated if additional correlation models become available to be tested. In general, it is reasonable that a correlation model altering the original covariance matrix as little as possible while improving the stability of the χ^2 leads to results similar to those obtained with the regularisation procedure. That being said, the shift in central value remains so small that it would be hard to make any conclusions based on its statistical significance.

On the other hand, the fact that using correlation models does not lead to a better χ^2 is a consequence of the fact that the condition number of the corresponding experimental covariance matrix is almost unaltered, as we have explicitly checked. We therefore conclude that the available correlation models are not enough to lead to a stable experimental

Table 3 The number of data points, N_{dat} , and the χ^2 per data point, χ^2/N_{dat} , for the NNPDF4.0 NNLO baseline fit, for each of the fits performed with a different correlation model (see text for details), and for the fit to the NNPDF4.0 data set regularised with $\delta^{-1} = 4$

Data set	N_{dat}	NNPDF4.0	χ^2/N_{dat}				$\delta^{-1} = 4$
			STRONG	WEAK	ATLAS	ATLAS+WEAK	
Deep-inelastic scattering	3089	1.12	1.12	1.12	1.12	1.12	1.11
Fixed-target Drell–Yan	195	0.98	1.00	0.99	0.99	0.99	0.97
Tevatron Drell–Yan	65	1.11	1.10	1.09	1.09	1.10	0.93
ATLAS total	679	1.24	1.24	1.24	1.23	1.24	1.04
W^+ +jet 8 TeV	15	0.79	0.78	0.79	0.79	0.79	0.79
W^- +jet 8 TeV	15	1.49	1.49	1.49	1.49	1.50	1.38
$t\bar{t}$ ℓ +jets y_t 8 TeV	4	3.28	3.14	3.06	3.04	3.16	2.98
$t\bar{t}$ ℓ +jets $y_{t\bar{t}}$ 8 TeV	4	3.83	3.58	3.57	3.54	3.65	3.40
jets R=0.6 8 TeV	171	0.68	0.69	0.68	0.68	0.68	0.69
dijets R=0.6 7 TeV	90	2.14	2.16	2.15	2.14	2.15	1.10
CMS total	474	1.31	1.31	1.31	1.31	1.30	1.21
LHCb total	116	1.55	1.56	1.54	1.55	1.55	1.56
Total	4618	1.16	1.16	1.16	1.16	1.16	1.11




Fig. 6 The PDFs obtained by fitting the NNPDF4.0 data set with correlation models provided by the experiment for a subset of measurements (see text for details). From top to bottom, left to right, we show the anti-up, anti-down, charm and gluon PDFs at a scale $Q = 100$ GeV. PDFs are compared to the NNPDF4.0 baseline parton set, and nor-

malised to its central value. Also shown are the PDFs obtained in a fit to the NNPDF4.0 data set regularised with $\delta^{-1} = 4$. For all PDFs but the NNPDF4.0 baseline, we show only the central value. Otherwise the uncertainty corresponds to the 68% confidence interval. All PDFs are accurate to NNLO

covariance matrix and χ^2 . In light of these considerations, we find that our regularisation procedure can possibly be used as a useful diagnosis tool to inform and validate correlation models not only at the level of PDF fits, but also at the level of the corresponding experimental analyses.

5 Conclusions

In this paper we have shown how an (even slightly) inaccurate determination of bin-by-bin correlations in the uncertainties of experimental measurements may make the χ^2 statistic fluctuate substantially, by more than one standard deviation. This problem is particularly relevant when dealing with high-precision measurements, in which the largest fraction of the uncertainty is correlated. This is the case for current and future LHC measurements that are routinely confronted with theoretical predictions by means of statistical inference. Because the χ^2 is routinely utilised as a figure of merit in these analyses, instabilities in its computation can make the interpretation of the results unreliable.

We have formulated the problem rigorously, by deriving a stability criterion for the acceptable fluctuations of the uncertainties on data correlations. The criterion ensures that the expectation value of the χ^2 does not overestimate its true value by an amount larger than its statistical fluctuation. To this aim, the criterion defines a bound on the singular values of the correlated part of the matrix of uncertainties. Building upon this criterion, we have then devised a regularisation procedure, whereby instabilities in the correlations of experimental uncertainties are removed with minimal information and without loss of generality. The idea is to clip the singular values of the correlated part of the matrix of uncertainties to a constant δ , whenever these are smaller than that, while leaving the rest of the singular vectors unchanged. This way, directions that do not contribute to instability are not affected and the alteration to the original matrix is minimal.

The key assumptions underlying the regularisation procedure are that correlations of experimental uncertainties across data points are determined much less precisely than the uncertainties on each data point, and that the prevalent source of inaccuracy on correlations concentrates on a subset of data points and originates from a small number of correlated uncertainties. The regularisation procedure leads to a covariance matrix that is more stable than the original one, when used to compute the χ^2 , is compatible with it within the precision with which it is determined, and does not lead to a reduction of the total uncertainty.

We have demonstrated how the regularisation procedure works in a toy model, and in a particular problem relevant to LHC precision physics that relies on the evaluation of the χ^2 as a figure of merit: PDF determination. Specifically, we have considered the NNPDF4.0 determination [9], which is

based on the widest data set to date. We have shown how the regularisation procedure can be utilised as a diagnosis tool to characterise the data set, in particular to single out those measurements for which an inaccurate estimation of experimental correlations may significantly affect their χ^2 . We have also studied how PDFs change if the nominal data set is replaced by a suitably regularised data set, and how these changes depend on the regularisation parameter δ . To this purpose, we have repeated the NNPDF4.0 baseline fit, now utilising a data set regularised with $\delta^{-1} = 1, 2, 3, 4, 5, 7$.

We have found that the χ^2 of some LHC data sets can be indeed significantly affected by inaccuracies in the determination of the correlations of their uncertainties. These inaccuracies can be reasonably regularised by choosing $\delta^{-1} = 4$. This value sets the precision with which uncertainties and correlations are known to less than 5% and 0.05, respectively. We have demonstrated that, by regularising the NNPDF4.0 data set with $\delta^{-1} = 4$, the global χ^2 is smaller than that of the baseline NNPDF4.0 determination by about 2.4σ . This means that it is only 5.3σ away from the unity expectation (instead of 7.7σ in the baseline NNPDF4.0 determination). At the same time, PDFs remain unaltered. These results highlight the fact that the nominal χ^2 determined in [9] is likely to be spuriously inflated by inaccuracies in the estimation of the experimental correlations in the LHC data.

Finally, we have studied how the regularisation procedure performs in comparison to correlation models provided with the measurements in the few cases in which these are available. We have found that our regularisation procedure captures the same qualitative shifts on PDF central values as experimental correlation models, once one or the others are used to determine the PDFs. Whether this is a coincidental feature, limited to the correlation models considered, or a more general one, could only be investigated if additional correlation models become available to be tested. On the other hand, using correlation models does not lead to a better χ^2 (or to a decrease in the condition number of the corresponding experimental covariance matrix). We therefore conclude that the available correlation models are not enough to lead to a stable χ^2 . In light of these considerations, we find that our regularisation procedure can possibly be used as a useful diagnosis tool to inform and validate correlation models not only at the level of PDF fits, but also at the level of the corresponding experimental analyses.

Our regularisation procedure is made publicly available as part of the NNPDF software [10]. The PDF sets discussed in this paper are available, in the LHAPDF format [54], from the NNPDF web page [55].

Acknowledgements We thank our colleagues of the NNPDF Collaboration for useful discussions, in particular Alessandro Candido and Christopher Schwan for a careful reading of the manuscript. Z. K. is supported by the European Research Council under the European Union's Horizon 2020 research and innovation Programme (grant agreement

n.950246). E. R. N. is supported by the U.K. Science and Technology Facility Council (STFC) grant ST/P000630/1. M. W. was supported by the STFC grant ST/R504737/1.

Data Availability Statement This manuscript has associated data in a data repository. [Authors' comment: The NNPDF software is publicly available from the github repository <https://github.com/NNPDF>. The PDF sets are publicly available from the NNPDF web page <https://nnpdf.mi.infn.it/nnpdf4-0-regularisation-study/>].

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

Funded by SCOAP³. SCOAP³ supports the goals of the International Year of Basic Sciences for Sustainable Development.

A Glossary

In this Appendix we recall the linear algebra definitions and the notation that are used throughout the paper. Let A be a $N \times M$ real matrix, $A \in \mathbb{R}^{N \times M}$, such as the error matrix introduced in Sect. 2. Then the following statements hold.

Singular value decomposition The matrix A admits a singular value decomposition of the form

$$A = USV^t, \tag{41}$$

where U and V are $N \times N$ and $M \times M$ orthogonal matrices, respectively, and S is a $N \times M$ diagonal matrix with non negative real numbers in the principal diagonal. The diagonal entries of S , s_i , $i = 1, \dots, \min(N, M)$, are the singular values of A , ordered by decreasing size, with $s_{\max} = s_1$ being the largest singular value.

Rank The rank of the matrix A , denoted as $rank(A)$, is the number of strictly positive singular values. The matrix A is full rank if its rank is $\min(N, M)$.

Pseudoinverse The pseudoinverse of the matrix A is

$$A^+ = VS^+U^t, \tag{42}$$

where S^+ is a $M \times N$ diagonal matrix with entries

$$S_{ii}^+ = \begin{cases} 1/s_i & s_i \neq 0 \\ 0 & s_i = 0 \end{cases}, \quad i = 1, \dots, \min(N, M). \tag{43}$$

Right inverse The right inverse of the matrix A (which we denote with the same notation used for the pseudoinverse) is the $M \times N$ matrix A^+ such that

$$AA^+ = I_{N \times N}. \tag{44}$$

Euclidean norm The Euclidean (or vector-induced L^2) norm of the matrix A is

$$\|A\|_2 = \max_{\{\mathbf{x} \in \mathbb{R}^M : \|\mathbf{x}\|=1\}} \|A\mathbf{x}\| = s_{\max}. \tag{45}$$

Frobenius norm The Frobenius norm of the matrix A is

$$\|A\|_F = \sqrt{\sum_i \sum_j A_{ij}^2} = \sqrt{tr(AA^t)} = \sqrt{\sum_i^{rank A} s_i^2}. \tag{46}$$

In Sects. 2–3 we always indicate with a subindex whether we are referring to the Euclidean or the Frobenius norm of the matrix A . By noting that $s_i \leq s_{\max}$, it follows from Eqs. (45) and (46) that

$$\|A\|_2 \leq \|A\|_F \leq \sqrt{rank(A)} \|A\|_2. \tag{47}$$

B Proof of Eq. (14)

In this Appendix, we provide a proof of Eq. (14). Let X and Y be arbitrary matrices of dimensions $M \times N$ and $N \times M$, respectively. We first show that

$$\|XY\|_F \leq \|X\|_F \|Y\|_2. \tag{48}$$

We write the respective singular value decompositions, $X = U_X S_X V_X^t$ and $Y = U_Y S_Y V_Y^t$, and define

$$W = V_X^t U_Y, \tag{49}$$

which is orthogonal because both factors are. Then

$$\|XY\|_F^2 = \|S_X W S_Y\|_F^2. \tag{50}$$

Hence

$$\|XY\|_F^2 = \sum_i \sum_j (S_{X(ii)} W_{(ij)} S_{Y(jj)})^2, \tag{51}$$

where the sums are implied to run over the relevant matrix dimensions M and N . We bind the singular values $S_{Y(jj)}$ by their maximum, the Euclidean norm of Y . That is

$$\|XY\|_F^2 \leq \sum_i \left(S_{X(ii)}^2 \sum_j W_{(ij)}^2 \|Y\|_2^2 \right). \tag{52}$$

Because W is orthogonal, it follows that

$$\|XY\|_F^2 \leq \left(\sum_i S_{X(ii)}^2 \right) \|Y\|_2^2 = \|X\|_F^2 \|Y\|_2^2. \tag{53}$$

The converse relation

$$\|XY\|_F \leq \|X\|_2 \|Y\|_F \tag{54}$$

can be proven analogously. Hence we obtain the result, Eq. (14),

$$\|XY\|_F \leq \min(\|X\|_F \|Y\|_2, \|X\|_2 \|Y\|_F). \tag{55}$$

References

1. M. Mangano, LHC at 10: the physics legacy. CERN Cour. **60**, 2–40 (2020). [arXiv:2003.05976](#)
2. P. Azzi et al., Report from Working Group 1: Standard Model Physics at the HL-LHC and HE-LHC. CERN Yellow Rep. Monogr. **7**, 1–220 (2019)
3. J.J. Ethier, E.R. Nocera, Parton distributions in nucleons and nuclei. Annu. Rev. Nucl. Part. Sci. **70**, 43–76 (2020). [arXiv:2001.07722](#)
4. X. Cid Vidal et al., Report from Working Group 3: Beyond the Standard Model physics at the HL-LHC and HE-LHC. CERN Yellow Rep. Monogr. **7**, 585–865 (2019). [arXiv:1812.07831](#)
5. I. Brivio, M. Trott, The standard model as an effective field theory. Phys. Rep. **793**, 1–98 (2019). [arXiv:1706.08945](#)
6. M. Cepeda et al., Report from Working Group 2: Higgs Physics at the HL-LHC and HE-LHC. CERN Yellow Rep. Monogr. **7**, 221–584 (2019). [arXiv:1902.00134](#)
7. G. Heinrich, Collider physics at the precision Frontier. Phys. Rep. **922**, 1–69 (2021). [arXiv:2009.00516](#)
8. P. Sinervo, Definition and treatment of systematic uncertainties in high energy physics and astrophysics, Statistical Problems in Particle Physics, Astrophysics, and Cosmology (2003)
9. NNPDF Collaboration, R.D. Ball et al., The path to proton structure at 1% accuracy. Eur. Phys. J. C **82**(5), 428 (2022). [arXiv:2109.02653](#)
10. NNPDF Collaboration, R.D. Ball et al., An open-source machine learning framework for global analyses of parton distributions. Eur. Phys. J. C **81**(10), 958 (2021). [arXiv:2109.02671](#)
11. E. Maguire, L. Heinrich, G. Watt, HEPData: a repository for high energy physics data. J. Phys. Conf. Ser. **898**, 10–102006 (2017). [arXiv:1704.05473](#)
12. ATLAS Collaboration, G. Aad et al., Measurement of the low-mass Drell–Yan differential cross section at $\sqrt{s} = 7$ TeV using the ATLAS detector. JHEP **06**, 112 (2014). [arXiv:1404.1212](#)
13. ATLAS Collaboration, G. Aad et al., Measurement of the high-mass Drell–Yan differential cross-section in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector. Phys. Lett. B **725**, 223–242 (2013). [arXiv:1305.4192](#)
14. ATLAS Collaboration, G. Aad et al., Measurement of the inclusive W^\pm and Z/γ cross sections in the electron and muon decay channels in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector. Phys. Rev. D **85**, 072004 (2012). [arXiv:1109.5141](#)
15. ATLAS Collaboration, M. Aaboud et al., Precision measurement and interpretation of inclusive W^+ , W^- and Z/γ^* production cross sections with the ATLAS detector. Eur. Phys. J. C **77**(6), 367 (2017). [arXiv:1612.03016](#)
16. ATLAS Collaboration, M. Aaboud et al., Measurement of the Drell–Yan triple-differential cross section in pp collisions at $\sqrt{s} = 8$ TeV. JHEP **12**, 059 (2017). [arXiv:1710.05167](#)
17. ATLAS Collaboration, G. Aad et al., Measurement of the double-differential high-mass Drell–Yan cross section in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. JHEP **08**, 009 (2016). [arXiv:1606.01736](#)
18. ATLAS Collaboration, G. Aad et al., Measurement of W^\pm and Z-boson production cross sections in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. Phys. Lett. B **759**, 601–621 (2016). [arXiv:1603.09222](#)
19. ATLAS Collaboration, M. Aaboud et al., Measurement of differential cross sections and W^+/W^- cross-section ratios for W boson production in association with jets at $\sqrt{s} = 8$ TeV with the ATLAS detector. JHEP **05**, 077 (2018) [Erratum: JHEP **10**, 048 (2020)]. [arXiv:1711.03296](#)
20. ATLAS Collaboration, G. Aad et al., Measurement of the transverse momentum and ϕ_η^* distributions of Drell–Yan lepton pairs in proton–proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. Eur. Phys. J. C **76**(5), 291 (2016). [arXiv:1512.02192](#)
21. ATLAS Collaboration, G. Aad et al., Measurements of top-quark pair differential cross-sections in the lepton+jets channel in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector. Eur. Phys. J. C **76**(10), 538 (2016). [arXiv:1511.04716](#)
22. ATLAS Collaboration, M. Aaboud et al., Measurement of top quark pair differential cross-sections in the dilepton channel in pp collisions at $\sqrt{s} = 7$ and 8 TeV with ATLAS. Phys. Rev. D **94**(9), 092003 (2016) [Addendum: Phys. Rev. D **101**, 119901 (2020)]. [arXiv:1607.07281](#)
23. ATLAS Collaboration, M. Aaboud et al., Measurement of the inclusive jet cross-sections in proton–proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. JHEP **09**, 020 (2017). [arXiv:1706.03192](#)
24. ATLAS Collaboration, G. Aad et al., Measurement of dijet cross sections in pp collisions at 7 TeV centre-of-mass energy using the ATLAS detector. JHEP **05**, 059 (2014). [arXiv:1312.3524](#)
25. ATLAS Collaboration, M. Aaboud et al., Measurement of the cross section for inclusive isolated-photon production in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector. Phys. Lett. B **770**, 473–493 (2017). [arXiv:1701.06882](#)
26. ATLAS Collaboration, G. Aad et al., Comprehensive measurements of t-channel single top-quark production cross sections at $\sqrt{s} = 7$ TeV with the ATLAS detector. Phys. Rev. D **90**(11), 112006 (2014). [arXiv:1406.7844](#)
27. ATLAS Collaboration, M. Aaboud et al., Fiducial, total and differential cross-section measurements of t-channel single top-quark production in pp collisions at 8 TeV using data collected by the ATLAS detector. Eur. Phys. J. C **77**(8), 531 (2017). [arXiv:1702.02859](#)
28. CMS Collaboration, S. Chatrchyan et al., Measurement of the electron charge asymmetry in inclusive W production in pp collisions at $\sqrt{s} = 7$ TeV. Phys. Rev. Lett. **109**, 111806 (2012). [arXiv:1206.2598](#)
29. CMS Collaboration, S. Chatrchyan et al., Measurement of the muon charge asymmetry in inclusive $pp \rightarrow W + X$ production at $\sqrt{s} = 7$ TeV and an improved determination of light parton distribution functions. Phys. Rev. D **90**(3), 032004 (2014). [arXiv:1312.6283](#)
30. CMS Collaboration, S. Chatrchyan et al., Measurement of the differential and double-differential Drell–Yan cross sections in proton–proton collisions at $\sqrt{s} = 7$ TeV. JHEP **12**, 030 (2013). [arXiv:1310.7291](#)
31. CMS Collaboration, V. Khachatryan et al., Measurement of the differential cross section and charge asymmetry for inclusive $pp \rightarrow W^\pm + X$ production at $\sqrt{s} = 8$ TeV. Eur. Phys. J. C **76**(8), 469 (2016). [arXiv:1603.01803](#)
32. CMS Collaboration, V. Khachatryan et al., Measurement of the Z boson differential cross section in transverse momentum and rapidity in proton–proton collisions at 8 TeV. Phys. Lett. B **749**, 187–209 (2015). [arXiv:1504.03511](#)
33. CMS Collaboration, S. Chatrchyan et al., Measurements of differential jet cross sections in proton–proton collisions at $\sqrt{s} = 7$ TeV with the CMS detector. Phys. Rev. D **87**(11), 112002 (2013) [Erratum: Phys. Rev. D **87**, 119902 (2013)]. [arXiv:1212.6660](#)
34. CMS Collaboration, V. Khachatryan et al., Measurement and QCD analysis of double-differential inclusive jet cross sections in pp collisions at $\sqrt{s} = 8$ TeV and cross section ratios to 2.76 and 7 TeV. JHEP **03**, 156 (2017). [arXiv:1609.05331](#)
35. CMS Collaboration, V. Khachatryan et al., Measurement of the differential cross section for top quark pair production in pp collisions at $\sqrt{s} = 8$ TeV. Eur. Phys. J. C **75**(11), 542 (2015). [arXiv:1505.04480](#)
36. CMS Collaboration, A.M. Sirunyan et al., Measurement of double-differential cross sections for top quark pair production in pp collisions at $\sqrt{s} = 8$ TeV and impact on parton distribution functions. Eur. Phys. J. C **77**(7), 459 (2017). [arXiv:1703.01630](#)

37. CMS Collaboration, A.M. Sirunyan et al., Measurements of $t\bar{t}$ differential cross sections in proton–proton collisions at $\sqrt{s} = 13$ TeV using events containing two leptons. *JHEP* **02**, 149 (2019). [arXiv:1811.06625](#)
38. CMS Collaboration, A.M. Sirunyan et al., Measurement of differential cross sections for the production of top quark pairs and of additional jets in lepton+jets events from pp collisions at $\sqrt{s} = 13$ TeV. *Phys. Rev. D* **97**(11), 112003 (2018). [arXiv:1803.08856](#)
39. LHCb Collaboration, R. Aaij et al., Measurement of the cross-section for $Z \rightarrow e^+e^-$ production in pp collisions at $\sqrt{s} = 7$ TeV. *JHEP* **02**, 106 (2013). [arXiv:1212.4620](#)
40. LHCb Collaboration, R. Aaij et al., Measurement of the forward Z boson production cross-section in pp collisions at $\sqrt{s} = 7$ TeV. *JHEP* **08**, 039 (2015). [arXiv:1505.07024](#)
41. LHCb Collaboration, R. Aaij et al., Measurement of forward $Z \rightarrow e^+e^-$ production at $\sqrt{s} = 8$ TeV. *JHEP* **05**, 109 (2015). [arXiv:1503.00963](#)
42. LHCb Collaboration, R. Aaij et al., Measurement of forward W and Z boson production in pp collisions at $\sqrt{s} = 8$ TeV. *JHEP* **01**, 155 (2016). [arXiv:1511.08039](#)
43. LHCb Collaboration, R. Aaij et al., Measurement of the forward Z boson production cross-section in pp collisions at $\sqrt{s} = 13$ TeV. *JHEP* **09**, 136 (2016). [arXiv:1607.06495](#)
44. T.-J. Hou et al., New CTEQ global analysis of quantum chromodynamics with high-precision data from the LHC. *Phys. Rev. D* **103**(1), 014013 (2021). [arXiv:1912.10053](#)
45. S. Bailey, T. Cridge, L.A. Harland-Lang, A.D. Martin, R.S. Thorne, Parton distributions from LHC, HERA, Tevatron and fixed target data: MSHT20 PDFs. *Eur. Phys. J. C* **81**(4), 341 (2021). [arXiv:2012.04684](#)
46. L.A. Harland-Lang, A.D. Martin, R.S. Thorne, The impact of LHC jet data on the MMHT PDF fit at NNLO. *Eur. Phys. J. C* **78**(3), 248 (2018). [arXiv:1711.05757](#)
47. S. Bailey, L. Harland-Lang, Differential top quark pair production at the LHC: challenges for PDF fits. *Eur. Phys. J. C* **80**(1), 60 (2020). [arXiv:1909.10541](#)
48. ATLAS Collaboration, G. Aad et al., Measurement of the inclusive jet cross-section in proton-proton collisions at $\sqrt{s} = 7$ TeV using 4.5 fb^{-1} of data with the ATLAS detector. *JHEP* **02**, 153 (2015) [Erratum: *JHEP* **09**, 141 (2015)]. [arXiv:1410.8857](#)
49. L.A. Harland-Lang, A.D. Martin, P. Motylinski, R.S. Thorne, Parton distributions in the LHC era: MMHT 2014 PDFs. *Eur. Phys. J. C* **75**(5), 204 (2015). [arXiv:1412.3989](#)
50. E.R. Nocera, M. Ubiali, Constraining the gluon PDF at large x with LHC data. *PoS DIS2017*, 008 (2018). [arXiv:1709.09690](#)
51. S. Amoroso et al., Les Houches 2019: Physics at TeV Colliders: Standard Model Working Group Report, in 11th Les Houches Workshop on Physics at TeV Colliders: PhysTeV Les Houches, 3 (2020). [arXiv:2003.01700](#)
52. NNPDF Collaboration, R.D. Ball et al., Parton distributions from high-precision collider data. *Eur. Phys. J. C* **77**(10), 663 (2017). [arXiv:1706.00428](#)
53. ATLAS Collaboration, G. Aad et al., Determination of the parton distribution functions of the proton using diverse ATLAS data from pp collisions at $\sqrt{s} = 7, 8$ and 13 TeV. *Eur. Phys. J. C* **82**(5), 438 (2022). [arXiv:2112.11266](#)
54. A. Buckley, J. Ferrando, S. Lloyd, K. Nordström, B. Page, M. Rüfenacht, M. Schönherr, G. Watt, LHAPDF6: parton density access in the LHC precision era. *Eur. Phys. J. C* **75**, 132 (2015). [arXiv:1412.7420](#)
55. <https://nnpdf.mi.infn.it/nnpdf4-0-regularisation-study/>