



Reweighting NNPDFs: The W lepton asymmetry

NNPDF Collaboration

Richard D. Ball ^{a,*}, Valerio Bertone ^b, Francesco Cerutti ^c,
Luigi Del Debbio ^a, Stefano Forte ^d, Alberto Guffanti ^b, José I. Latorre ^c,
Juan Rojo ^d, Maria Ubiali ^e

^a *Tait Institute, University of Edinburgh, JCMB, KB, Mayfield Rd, Edinburgh EH9 3JZ, Scotland, United Kingdom*

^b *Physikalisches Institut, Albert-Ludwigs-Universität Freiburg, Hermann-Herder-Straße 3,
D-79104 Freiburg i. B., Germany*

^c *Departament d'Estructura i Constituents de la Matèria, Universitat de Barcelona, Diagonal 647,
E-08028 Barcelona, Spain*

^d *Dipartimento di Fisica, Università di Milano and INFN, Sezione di Milano, Via Celoria 16, I-20133 Milano, Italy*

^e *Institut für Theoretische Teilchenphysik und Kosmologie, RWTH Aachen University, D-52056 Aachen, Germany*

Received 7 December 2010; received in revised form 9 March 2011; accepted 14 March 2011

Available online 23 March 2011

Abstract

We present a method for incorporating the information contained in new datasets into an existing set of parton distribution functions without the need for refitting. The method involves reweighting the ensemble of parton densities through the computation of the χ^2 to the new dataset. We explain how reweighting may be used to assess the impact of any new data or pseudodata on parton densities and thus on their predictions. We show that the method works by considering the addition of inclusive jet data to a DIS + DY fit, and comparing to the refitted distribution. We then use reweighting to determine the impact of recent high statistics lepton asymmetry data from the D0 experiment on the NNPDF2.0 parton set. We find that the D0 inclusive muon and electron data are perfectly compatible with the rest of the data included in the NNPDF2.0 analysis and impose additional constraints on the large- x d/u ratio. The more exclusive D0 electron datasets are however inconsistent both with the other datasets and among themselves, suggesting that here the experimental uncertainties have been underestimated.

Crown Copyright © 2011 Published by Elsevier B.V. All rights reserved.

* Corresponding author.

E-mail address: rdb@ph.ed.ac.uk (R.D. Ball).

1. Introduction

The determination of parton distribution functions (PDFs) and their uncertainties through global fits to datasets taken in deep inelastic and hadronic collision experiments, analyzed using perturbative QCD, is one of the key ingredients in the exploitation of future experiments, in particular at LHC. Of course such fits can only be as good as the data that goes into them, so whenever there is new data or new experiments, the fits have to be redone to take the new data into account. This process is cumbersome and time consuming, and can only be performed using the same software as in the previous fits, and thus by the fitting collaborations themselves.

In this paper we will show how, by using the ensembles of PDFs produced by the NNPDF Collaboration [1–4], anyone can determine the effect of new data on the PDFs quickly and easily: all that is required is a computation of the χ^2 to the new data for each PDF in the ensemble [5]. With this information one can determine the overall impact of the new data on PDFs, their consistency with the older data used in the fit, the effect the new data have on the shape and precision of individual PDFs, and thus their effect on observables such as benchmark cross-sections or predictions for new physics. The same approach can be used just as easily to estimate the effects of pseudodata from proposed experiments or machines on PDFs and thus on cross-sections.

The technique we use is based on statistical inference. In the NNPDF approach [1–4,6] we generate through a Monte Carlo procedure an ensemble of N PDF replicas, $\mathcal{E} = \{f_k, k = 1, \dots, N\}$, each fitted to a data replica generated according to the uncertainties and their correlations as measured by the experimental collaborations. Each PDF is parametrized by a highly redundant neural network in order to avoid parameterization bias which would otherwise spoil the procedure, and the stopping point of the fit of each replica is determined using cross-validation to prevent over-fitting. The final PDF ensemble then forms an accurate representation of the probability distribution of PDFs,¹ conditional on the input data and the particular assumptions (such as NLO QCD, a value of α_s , etc.) used in the fits.

Given an NNPDF ensemble one can evaluate any quantity or experimental observable $\mathcal{O}[f]$ depending on the PDFs (such as the PDF mean, the variance, PDF correlations, or indeed the mean, variance, etc., of any cross-section computed from them) by computing $\mathcal{O}[f]$ for each of the replicas, and averaging the results. This is because the integral in the space of functions is well approximated by the average over the ensemble \mathcal{E} , so that the mean value of $\mathcal{O}[f]$ given by

$$\langle \mathcal{O} \rangle = \int \mathcal{O}[f] \mathcal{P}(f) Df = \frac{1}{N} \sum_{k=1}^N \mathcal{O}[f_k]. \quad (1)$$

Each of the replicas f_k carries equal weight because they were generated through importance sampling: the replicas were fitted to a data replica generated according to the probability distribution of the experimental data, using a fitting procedure with no bias.

We can include the effects of a new independent dataset without performing a new fit if we instead reweight the old fit according to weights w_k , which assess the probability that each PDF replica f_k agrees with the new data. The reweighted ensemble then forms a representation of the probability distribution of PDFs conditional on both the old data and the new. The weights are computed straightforwardly by evaluating the χ^2 of the new data to each of the replicas. The

¹ Throughout this paper we will denote ‘parton distribution function’ by PDF, but write out ‘probability density function’ in full, in order to avoid any confusion: both are probability densities, but in very different spaces.

mean value of the observable $\mathcal{O}[f]$ taking account of the new data is then given by the weighted average

$$\langle \mathcal{O} \rangle_{\text{new}} = \int \mathcal{O}[f] \mathcal{P}_{\text{new}}(f) Df = \frac{1}{N} \sum_{k=1}^N w_k \mathcal{O}[f^{(k)}]. \quad (2)$$

The usefulness of this method is clear: it becomes possible to test the impact of a new dataset, or indeed the potential impact of MC data generated for a new experiment, quickly and simply without the need for a new fit (or indeed without considering explicitly any other datasets except the one under immediate consideration). This comes at a price: the effective number of replicas will be reduced, either because the new data prove to be very constraining (good), or because they are inconsistent with the old data (bad). We will provide a criterion to distinguish between these two cases. Of course if the new data are both consistent and precise, the effective number of replicas might be so reduced that a refit becomes necessary.

One of the advantages of the reweighting method is that it can be used to assess the impact on the global fit of observables for which no fast code is available, and thus which cannot be included without resorting to K -factor approximations. One such observable is the Tevatron W lepton charge asymmetry. Recent measurements [7–10] of this quantity have attracted a lot of attention, since sizable tension with other data in the global fit sensitive to the large- x d quark distribution, such as DIS deuterium structure function data, has been reported [11,12]. It is not clear from these studies whether this tension reflects an experimental problem of the recent Tevatron data, or whether the problems are with the DIS deuterium data, perhaps indicating the need for substantial nuclear corrections. With this motivation, armed with the statistical power of reweighting, we will here study the incorporation of the D0 W lepton charge asymmetry data on the NNPDF2.0 fit.

Reweighting is also important from a conceptual point of view. If more and more data are included in the fit through reweighting, the resulting PDFs become less and less dependent on the initial PDF. But PDFs obtained in this way then by construction satisfy the laws of statistical inference — for example, uncertainties will automatically behave upon inclusion of new data according to standard statistics.² Hence, checking that the results obtained by reweighting coincide with results obtained by simply including the new data in the global fit provides a highly nontrivial check on the consistency of the NNPDF global fitting procedure.

The paper is organized as follows. In the next section, we will explain how the weights can be computed, and give tests through which one can access quantitatively the impact of the new data and their consistency with the old data. A detailed proof of these results, with a full discussion of the subtleties, is given in Section 3: this is important because an earlier attempt to implement a reweighting procedure [5] used an expression for the weights which differs from our result (a detailed examination of the result of Ref. [5] is presented in Section 3.2).³ This section may be skipped by readers only interested in applying the new technique. In Section 4 we show how the method may be used in practice by applying it to inclusive jet data: since there are existing NNPDF sets with and without this data, this allows us to check that reweighting works. Then in Section 5 we illustrate the power of reweighting by using it to assess the impact of D0 W lepton

² Indeed, it was suggested in Ref. [13] that a PDF fit might be performed by including all data through reweighting of a first guess based on past experience.

³ A recent study by the LHCb Collaboration [14] using a reweighting technique to assess the impact of low p_t Drell–Yan pairs at the LHC on PDF determinations, also appears to use the incorrect formula derived in [5].

charge asymmetry data on the NNPDF2.0 PDFs. The results are particularly interesting because while the inclusive D0 asymmetry data is perfectly compatible with the NNPDF2.0 set and results in a moderate improvement in the determination of the medium and large- x d quark PDF, the more exclusive electron datasets turn out to be inconsistent both with other sets in the global analysis and among them.

2. Reweighting

2.1. The weights

We consider the situation where a set of experimental data have been used to construct a probability distribution for PDFs, $\mathcal{P}_{\text{old}}(f)$. This probability distribution is delivered as a finite ensemble of PDFs $\mathcal{E} = \{f^{(k)}, k = 1, \dots, N\}$. Any observable can be obtained performing averages over this ensemble as prescribed in Eq. (1), that is, equally weighting each PDF.

The problem we shall now address is how to update this probability distribution when new experimental data are available. There are two options: either we can construct a new probability distribution $\mathcal{P}_{\text{new}}(f)$ from scratch by including both old and new data in a new fit, or we can update the old fit by computing a weight, w_k , for each individual PDF f_k in the ensemble \mathcal{E} according to the rules of statistical inference. Then, $\mathcal{P}_{\text{new}}(f)$ is simply understood as an update (or reweighting) of the prior probability distribution $\mathcal{P}_{\text{old}}(f)$.

Both methods incorporate the same information, the old and the new data, and we will show below that when the weights are chosen correctly both techniques are statistically equivalent in the sense that when the number of replicas is sufficiently large they both give representations of the same probability distribution $\mathcal{P}_{\text{new}}(f)$. However to calculate the weights involves only knowledge of the new data: all the relevant information about the old data is already contained in $\mathcal{P}_{\text{old}}(f)$. It is thus substantially easier to implement, since no refitting is necessary.

To be specific, we consider a set of n new data that have not been included in the determination of the initial probability density distribution:

$$y = \{y_1, y_2, \dots, y_n\}.$$

Clearly any instance of such a set of data can be seen as a point y in an n -dimensional real space. The experimental uncertainties are summarized by the $n \times n$ experimental covariance matrix σ_{ij} , which includes a term that incorporates the overall normalization uncertainty [15], but reduces to a diagonal matrix in cases when experimentalists do not provide the correlated systematic uncertainties. We assume throughout that these new data are statistically independent of any of the data included in the original fit.

Using statistical inference we can update the initial probability density $\mathcal{P}_{\text{old}}(f)$ by taking into account the new data, thereby obtaining an improved probability density $\mathcal{P}_{\text{new}}(f)$. To do this we need to know the relative probabilities of the new data for different choices of PDF. Since the new data are assumed to have Gaussian errors, these probabilities will be directly proportional to the probability density of the χ^2 to the new data conditional on f :

$$\mathcal{P}(\chi^2|f) \propto (\chi^2(y, f))^{n/2-1} e^{-\frac{1}{2}\chi^2(y, f)}, \quad (3)$$

where if $y_i[f]$ is the value predicted for the data y_i using the PDF f ,

$$\chi^2(y, f) = \sum_{i,j=1}^n (y_i - y_i[f])\sigma_{ij}^{-1}(y_j - y_j[f]). \quad (4)$$

It follows from the statistical independence of the old and new data that by the law of multiplication of probabilities:

$$\mathcal{P}_{\text{new}}(f) = \mathcal{N}'_{\chi} \mathcal{P}(\chi^2|f) \mathcal{P}_{\text{old}}(f), \quad (5)$$

where \mathcal{N}'_{χ} is a normalization factor, independent of f .

Multiplying on both sides by some observable $\mathcal{O}[f]$ and integrating over the PDFs,

$$\begin{aligned} \langle \mathcal{O} \rangle_{\text{new}} &= \int \mathcal{O}[f] \mathcal{P}_{\text{new}}(f) Df \\ &= \mathcal{N}'_{\chi} \int \mathcal{O}[f] \mathcal{P}(\chi^2|f) \mathcal{P}_{\text{old}}(f) Df \\ &= \frac{1}{N} \sum_{k=1}^N \mathcal{N}'_{\chi} \mathcal{P}(\chi^2|f_k) \mathcal{O}[f_k], \end{aligned} \quad (6)$$

where in the last line we used Eq. (1).

We can thus sample the probability density $\mathcal{P}_{\text{new}}(f)$ using the N replicas f_k , but reweighted: in place of Eq. (1) we now have

$$\langle \mathcal{O} \rangle_{\text{new}} = \frac{1}{N} \sum_{k=1}^N w_k \mathcal{O}[f_k], \quad (7)$$

where

$$w_k = \mathcal{N}'_{\chi} \mathcal{P}(\chi^2|f_k) = \mathcal{N}'_{\chi} (\chi_k^2)^{n/2-1} e^{-\frac{1}{2}\chi_k^2}, \quad (8)$$

and $\chi_k^2 \equiv \chi^2(y, f_k)$ is evaluated using Eq. (4). The normalization factor \mathcal{N}'_{χ} is fixed by normalizing the new probability density $\mathcal{P}_{\text{new}}(f)$: taking the operator $\mathcal{O}[f]$ to be the unit operator, $\langle 1 \rangle_{\text{new}} = 1$, so from Eq. (7) this fixes $\sum_{k=1}^N w_k = N$, and thus using Eq. (8)

$$w_k = \frac{(\chi_k^2)^{n/2-1} e^{-\frac{1}{2}\chi_k^2}}{\frac{1}{N} \sum_{k=1}^N (\chi_k^2)^{n/2-1} e^{-\frac{1}{2}\chi_k^2}}. \quad (9)$$

The weights w_k , when divided by N , are then simply the probabilities of the replicas f_k , given the χ^2 to the new data.

Note that our formula (9) for the weights is different from the one derived in Ref. [5]. The reason for this is that the use of Bayes theorem for multi-dimensional probability densities is subtle, since without care one may fall foul of the Borel–Kolmogorov paradox (see for example Ref. [16]). A careful proof of the weights equation (9) using the elementary rules of statistical inference is given in Section 3.1: the subtle error in the argument used in Ref. [5] is explained in Section 3.2.

2.2. Measuring information loss and consistency

The original ensemble of replicas $\mathcal{E} = \{f^{(k)}, k = 1, \dots, N\}$ is constructed through importance sampling of the probability density $\mathcal{P}_{\text{old}}(f)$, and thus each replica has equal weight. It is maximally efficient, in the sense this is the best representation of the underlying density $\mathcal{P}_{\text{old}}(f)$ for a given number of replicas N : the only way to improve it is by increasing N . After reweighting, this will no longer be the case, since in fact the weights give the relative importance of the

different replicas, and the replicas with very small weights will become almost irrelevant in ensemble averages. The reweighted replicas will thus no longer be as efficient as the old: for a given N , the accuracy of the representation of the underlying distribution $\mathcal{P}_{\text{new}}(f)$ will be less than it would be in a new fit.

We can quantify this loss of efficiency by using the Shannon entropy to compute the effective number of replicas left after reweighting:

$$N_{\text{eff}} \equiv \exp \left\{ \frac{1}{N} \sum_{k=1}^N w_k \ln(N/w_k) \right\}. \quad (10)$$

Clearly $0 < N_{\text{eff}} < N$: the reweighted fit has the same accuracy as a refit with N_{eff} replicas. Thus if N_{eff} becomes too low, the reweighting procedure will no longer be reliable, either because the new data contain a lot of information on the PDFs, necessitating a full refitting with more replicas, or because the new data are inconsistent with the old.

These two cases can be distinguished by examining the χ^2 profile of the new data: if in the reweighted fit there are very few replicas with a χ^2 per data point of order unity, the errors in the new dataset have probably been underestimated. This profile may be easily evaluated:

$$\mathcal{P}(\chi^2) = \frac{1}{N} \sum_k w_k, \quad (11)$$

where the sum is over all replicas k such that $\chi_k^2 \in [\chi^2, \chi^2 + d\chi^2]$.

Alternatively, we consider inconsistent data as data whose errors have been underestimated. If we rescale the covariance matrix of the data by a factor α , we can use inverse probability to calculate the probability density for the rescaling parameter α :

$$\mathcal{P}(\alpha) \propto \frac{1}{\alpha} \sum_{k=1}^N w_k w_k(\alpha). \quad (12)$$

Here $w_k(\alpha)$ are the weights equation (9) evaluated by replacing χ_k^2 with χ_k^2/α (and are thus proportional to the probability of f_k given the new data with rescaled errors): averaging them in the reweighted fit thus gives the probability density for α . If this probability density peaks close to one the new data are consistent, while if it peaks far above one, then it is likely that the errors in the data have been underestimated.

If the new data are reasonably consistent, we can assess whether they should be included in a new fit by calculating the χ^2 distribution of the dataset that would be used in the new fit (i.e. all the old data plus the new data), using the reweighting procedure as in Eq. (11). Comparison to the old χ^2 distribution then tells us whether the new data would improve the fit: if so the peak should shift a little towards one, with a slight narrowing due to the increase in the total number of data points. This may be quantified by comparing the area under the curves in a given range.

3. Statistical inference

3.1. Proof of the weight formula

Here we give a careful derivation of the rules for reweighting presented in the previous section. Some readers may consider skipping this section and simply use the practical prescription as given in Eq. (9). Our argument is based on the standard use of statistical inference. However

some of the details are subtle, since we need to use probability densities in multi-dimensional spaces, and thus need to take care with limits.

By the probability $P(f)$ for the PDF f what we actually mean is the probability $P(f|K)$, where K denotes all the data used in the determination and their associated errors, the values of parameters such as α_s and heavy quark masses used in the computation of the data expected from the PDF, and finally also the theoretical framework used (for example NLO QCD). If we then wish to extend the dataset by including new data y , the new probability $P_{\text{new}}(f)$ is then $P(f|yK)$: besides K we now also assume the new data y .

The new probability is then determined from the old probability using the sampling distribution $P(y|fK)$ and multiplicative rule for probabilities (often known as Bayes theorem):

$$P(AB|C) = P(A|BC)P(B|C) = P(B|AC)P(A|C), \quad (13)$$

where $P(A|C)$ is the probabilities of A given C , etc., and AB denotes A and B . Naively applying this result in the present case we have

$$P(f|yK)P(y|K) = P(y|fK)P(f|K), \quad (14)$$

whence (replacing $P(f|K)$ with $\mathcal{P}(f|K) Df$, $P(f|yK)$ with $\mathcal{P}(f|yK) Df$)

$$\mathcal{P}(f|yK) = P(y|fK)\mathcal{P}(f|K)/P(y|K). \quad (15)$$

Note that $P(y|K)$ does not depend on the PDF f , and can thus be determined simply by insisting that $\mathcal{P}(f|yK)$ is properly normalized: we then find

$$P(y|K) = \int P(y|fK)\mathcal{P}(f|K) Df, \quad (16)$$

so

$$\mathcal{P}(f|yK) = P(y|fK)\mathcal{P}(f|K) / \int P(y|fK)\mathcal{P}(f|K) Df, \quad (17)$$

where everything on the right-hand side is now known.

This argument would work without problems if the data y could only take discrete values. The difficulty in the present case is that our data are continuous, so rather than the probability $P(y|fK)$ we have to work with a multi-dimensional probability density $\mathcal{P}(y|fK) d^n y$, in a limit in which the volume element $d^n y$ goes to zero. Of course in this limit the probabilities $P(y|fK)$ and $P(y|K)$ also go to zero, and we find a ratio of two zeros in Eq. (15). The conditional probability $P(f|yK)$ is then only well defined if we specify carefully the way in which the limit is to be taken: probabilities conditional on sets of measure zero are ambiguous. Failure to specify the limiting process can result in contradictions (the Borel–Kolmogorov paradox [16]).

Consider then the probability density for the data y . Assuming that the new experiments are not correlated with any of the experiments used in the determination of the initial probability density, the probability density of y is then given by Eq. (16):

$$P(y|K) = \int \mathcal{P}(y|fK)\mathcal{P}(f|K) Df = \frac{1}{N} \sum_{k=1}^N \mathcal{P}(y|f_k K), \quad (18)$$

where in the second step we used Eq. (1). The density $\mathcal{P}(y|fK)$ gives the probability that the new data lie in an infinitesimal volume $d^n y$ centred at y in the space of possible data given a particular choice of PDF f : it is often called the sampling distribution or (when considered as

a function of f) the likelihood function. Assuming that the uncertainties in the data are purely Gaussian,

$$\mathcal{P}(y|fK) d^n y = (2\pi)^{-n/2} (\det \sigma_{ij})^{-1/2} e^{-\frac{1}{2}\chi^2(y,f)} d^n y, \quad (19)$$

where $\chi^2(y, f)$ is calculated using Eq. (4) (and of course using the assumptions K in the computation of the predictions $y_i[f]$).

Since to compute $\mathcal{P}(y|fK)$ it is sufficient to compute $\chi^2(y, f)$, it is sufficient for our purposes to consider the probability density for the χ^2 to the new dataset:

$$\mathcal{P}(\chi^2|fK) d\chi^2 = 2^{-n/2} (\Gamma(n/2))^{-1} (\chi^2(y, f))^{n/2-1} e^{-\frac{1}{2}\chi^2(y,f)} d\chi^2. \quad (20)$$

This may be readily derived from Eq. (19) by diagonalizing the covariance matrix and rescaling the data to a set $\{Y_i\}$ of independent Gaussian variables each with unit variance. Then $d^n y = (\det \sigma_{ij})^{1/2} d^n Y$, and $\chi^2 = \sum_{i=1}^n Y_i^2$. Choosing n -dimensional spherical co-ordinates in the space of data (with $\sqrt{\chi^2}$ as the radial co-ordinate, and thus $y = y[f]$ as the origin), we may write $d^n Y = A_n \frac{1}{2} (\chi^2)^{n/2-1} d\chi^2 d^{n-1} \Omega$, where $d^{n-1} \Omega$ is the measure on the sphere and $A_n = 2\pi^{n/2} (\Gamma(n/2))^{-1}$ is the area of the unit sphere in n dimensions. The probability (19) may thus be written

$$\begin{aligned} \mathcal{P}(y|fK) d^n y &= (2\pi)^{-n/2} e^{-\frac{1}{2}\chi^2(y,f)} d^n Y \\ &= 2^{-n/2} (\Gamma(n/2))^{-1} (\chi^2(y, f))^{n/2-1} e^{-\frac{1}{2}\chi^2(y,f)} d\chi^2 d^{n-1} \Omega, \end{aligned} \quad (21)$$

in agreement with Eq. (20) provided

$$\mathcal{P}(y|fK) d^n y = \mathcal{P}(\chi^2|fK) d\chi^2 d^{n-1} \Omega. \quad (22)$$

Again the probability density $\mathcal{P}(\chi^2|K)$ for the χ^2 of the new dataset is obtained by averaging over replicas:

$$\mathcal{P}(\chi^2|K) = \int \mathcal{P}(\chi^2|fK) \mathcal{P}(f|K) Df = \frac{1}{N} \sum_{k=1}^N \mathcal{P}(\chi^2|f_k K); \quad (23)$$

so combining Eqs. (18), (22), and (23)

$$\mathcal{P}(y|K) d^n y = \frac{1}{N} \sum_{k=1}^N \mathcal{P}(\chi^2|f_k K) d\chi^2 d^{n-1} \Omega = \mathcal{P}(\chi^2|K) d\chi^2 d^{n-1} \Omega, \quad (24)$$

since both the volume factor $d^{n-1} \Omega$ and the interval $d\chi^2$ are independent of the choice of replica, and may thus be taken out of the sum.

The advantage of using $\mathcal{P}(\chi^2|fK)$ instead of $\mathcal{P}(y|fK)$ when evaluating Eq. (15) is that $\mathcal{P}(\chi^2|fK)$ is only a one-dimensional density, so taking the limit in which the volume element goes to zero is straightforward and unambiguous. We may write Eq. (15) as

$$\mathcal{P}(f|\chi^2 K) Df \mathcal{P}(\chi^2|K) d\chi^2 = \mathcal{P}(\chi^2|fK) d\chi^2 \mathcal{P}(f|K) Df. \quad (25)$$

The marginalization equation (23) follows directly on integration over f , since if $\mathcal{P}(f|\chi^2)$ is correctly normalized, $\int \mathcal{P}(f|\chi^2 K) Df = 1$. Now, cancelling the $d\chi^2$ from either side of Eq. (25) (since this is just a pre-assigned interval),

$$\mathcal{P}(f|\chi^2 K) Df = \frac{\mathcal{P}(\chi^2|fK)}{\mathcal{P}(\chi^2|K)} \mathcal{P}(f|K) Df. \quad (26)$$

Multiplying on both sides by some observable $\mathcal{O}[f]$ and integrating over the PDFs,

$$\begin{aligned} \langle \mathcal{O} \rangle_{\text{new}} &= \int \mathcal{O}[f] \mathcal{P}(f | \chi^2 K) Df \\ &= \int \mathcal{O}[f] \frac{\mathcal{P}(\chi^2 | f K)}{\mathcal{P}(\chi^2 | K)} \mathcal{P}(f | K) Df \\ &= \frac{1}{N} \sum_{k=1}^N \frac{\mathcal{P}(\chi^2 | f_k K)}{\mathcal{P}(\chi^2 | K)} \mathcal{O}[f_k], \end{aligned} \quad (27)$$

where in the last line we used Eq. (1). This corresponds to the reweighting Eq. (7) with weights

$$w_k = \frac{\mathcal{P}(\chi^2 | f_k K)}{\mathcal{P}(\chi^2 | K)}. \quad (28)$$

Combining Eq. (28) with Eqs. (20) and (23), we obtain Eq. (9).

Note that a further application of Bayes' theorem to Eq. (28) gives the alternative form

$$w_k = \frac{P(f_k | \chi^2 K)}{P(f_k | K)} = N P(f_k | \chi^2 K), \quad (29)$$

since because the replicas are uniformly distributed, $P(f_k | K) = 1/N$. Thus w_k/N is the probability of replica f_k given the χ^2 to the new data.

3.2. The naive prescription

It is instructive to also derive the weights working directly with the probability density $\mathcal{P}(y | f K)$: using Bayes' theorem we may write instead of Eq. (25)

$$\mathcal{P}(f | y K) Df \mathcal{P}(y | K) d^n y = \mathcal{P}(y | f K) d^n y \mathcal{P}(f | K) Df. \quad (30)$$

Again, the marginalization equation (18) follows directly from the requirement that $\mathcal{P}(f | y K)$ be normalized, i.e. that $\int \mathcal{P}(f | y K) Df = 1$.

Naively cancelling the volume factor $d^n y$ from either side, and pursuing the same argument as before yields:

$$\begin{aligned} \langle \mathcal{O} \rangle_{\text{new}}^G &= \int \mathcal{O}[f] \mathcal{P}(f | y K) Df \\ &= \int \mathcal{O}[f] \frac{\mathcal{P}(y | f K)}{\mathcal{P}(y | K)} \mathcal{P}(f | K) Df \\ &= \frac{1}{N} \sum_{k=1}^N \frac{\mathcal{P}(y | f_k K)}{\mathcal{P}(y | K)} \mathcal{O}[f_k]. \end{aligned} \quad (31)$$

This would then lead to the conclusion of Giele and Keller [5] that the weights are proportional to $\mathcal{P}(y | f_k K) / \mathcal{P}(y | K)$, and thus (using Eq. (19)) are given by

$$w_k^G = \frac{e^{-\frac{1}{2}\chi_k^2}}{\frac{1}{N} \sum_{k=1}^N e^{-\frac{1}{2}\chi_k^2}}. \quad (32)$$

This result is clearly different from our previous result (28). We see explicitly that the densities $\mathcal{P}(f | \chi^2 K)$ and $\mathcal{P}(f | y K)$ are not the same, despite the fact that when the data y take a given value, the χ^2 takes a corresponding value.

It is also clear that the Gaussian weights equation (32) must be incorrect: in the limit where the number of new (and consistent) data n becomes very large, χ_k^2 peaks around n , and only the very few replicas in the tail of the distribution ($\chi_k^2 \ll n$ is very unlikely) will survive. By contrast with the correct weights equation (28), replicas with $\chi_k^2 \sim n$ will dominate the fit, replicas with very small or very large χ^2 being suppressed.

The reason for the difference between the results (28) and (32) is the Borel–Kolmogorov paradox [16]: when dealing with multi-dimensional probability densities care must be taken with limits, since a conditional probability on a set of measure zero is not well defined. Here the limits used to derive $\mathcal{P}(f|\chi^2 K)$ and $\mathcal{P}(f|yK)$ are different, and thus so are the results. The correct result can only be obtained by taking the appropriate limit.

The probability density $\mathcal{P}(f|\chi^2 K)$ is defined as the probability density for f given that the χ^2 lies in the finite interval $[\chi^2, \chi^2 + d\chi^2]$, in the limit $d\chi^2 \rightarrow 0$. In this case the conditioning variable spans a one-dimensional manifold, and therefore there is no freedom in the choice of the limiting procedure. The definition of $\mathcal{P}(f|\chi^2 K)$ is unique, and thus the argument which leads to Eq. (28) unambiguous. However the probability density $\mathcal{P}(f|yK)$ is defined as the probability density for f given that y lies in some volume V_n , in the limit $V_n \rightarrow 0$. In a multi-dimensional space such as this, the conditional probability density $\mathcal{P}(f|yK)$ is ambiguous, since it depends on the way the volume element V_n is chosen, and then taken to zero. Different definitions correspond to different physical settings. In the argument which led us to Eq. (32), we implicitly assumed that V_n was the compact volume $d^n y$ centred on y , so that as $V_n \rightarrow 0$, the point y was uniquely selected. However there are many points in the space of data which have the same χ^2 , and thus the same effect on f . If we want to include all these points with equal weight when determining the conditional probability density of f given y , as we surely must, we need to sum over all the compact volumes $d^n y$ that build the $(n - 1)$ -dimensional level surfaces of $\chi^2(y, f)$ through the point y . Thus V_n is a thin shell with thickness $d\chi^2$, and hence its total volume is $V_n = A_n d\chi^2$. The limit $V_n \rightarrow 0$ is then taken by letting $d\chi^2 \rightarrow 0$. We should thus write Eq. (30) as (using Eqs. (22) and (24))

$$\mathcal{P}(f|yK) Df \mathcal{P}(\chi^2|K) A_n d\chi^2 = \mathcal{P}(\chi^2|fK) A_n d\chi^2 \mathcal{P}(f|K) Df. \tag{33}$$

Cancelling the volume factor $A_n d\chi^2$, since this is independent of f , this definition is the same as Eq. (26), and thus yields the correct weights equation (9) in the limit $V_n \rightarrow 0$.

3.3. Derivation of the consistency tests

Finally we consider the derivation of the two diagnostic results (11) and (12). The first is simply the ‘evidence’ (23), evaluated by binning in χ^2 . The second is more involved: using Bayes theorem

$$\mathcal{P}(\alpha|\chi^2, f_k, K) d\alpha = \frac{P(f_k|\alpha, \chi^2, K)}{P(f_k|\chi^2, K)} \mathcal{P}(\alpha|\chi^2, K) d\alpha. \tag{34}$$

Now the probability $P(f_k|\alpha, \chi^2, K)$ may be evaluated by noting that the effect of α is to rescale $\chi^2 \rightarrow \chi^2/\alpha$, and $P(f_k|\chi^2/\alpha, K)$ is then simply proportional to the weights $w_k(\alpha)$ evaluated according to Eq. (9) with χ^2 replaced by χ^2/α . The prior density $\mathcal{P}(\alpha|\chi^2, K)$ we assume is uniform in $\ln\alpha$, since α is a scale parameter (this ensures that the results are invariant under $\alpha \rightarrow 1/\alpha$). We thus find

$$\mathcal{P}(\alpha|\chi^2, f_k, K) = \frac{w_k(\alpha)}{\alpha \int d(\ln\alpha') w_k(\alpha')}, \quad (35)$$

where the overall normalization has been fixed by integrating over α . Then as usual

$$\begin{aligned} \mathcal{P}(\alpha|\chi^2, K) &= \int Df \mathcal{P}(\alpha|\chi^2, f, K) \mathcal{P}(f|\chi^2, K) \\ &= \frac{1}{N} \sum_{k=1}^N \frac{w_k w_k(\alpha)}{\alpha \int d(\ln\alpha') w_k(\alpha')}. \end{aligned} \quad (36)$$

It is easy to show by a change of variable that the integrals in the denominator are the same for all k , whence we find Eq. (12).

4. Validation: Inclusive jets

As a demonstration of the effectiveness of our reweighting procedure, we first apply it to a dataset that has already been included and studied in the NNPDF2.0 analysis [4]. We thus start with the fit obtained including only the DIS and Drell–Yan data, call this NNPDF2.0 (DIS + DYP), and then add the inclusive jet data from Tevatron Run II [17,18], which were included in the NNPDF2.0 analysis, through reweighting. The resulting reweighted fit can then be compared directly with the NNPDF2.0 fit, which includes the same DIS, Drell–Yan and Tevatron inclusive jet data. Given the consistency of the inclusive jet data with the DIS and Drell–Yan data demonstrated in Ref. [4], we expect the reweighted and refitted distributions to give results that are equivalent up to statistical fluctuations.

Note that from this section on we will slightly change the notation to make it more similar to that of previous NNPDF studies: N_{rep} will denote the number of replicas in the sample and N_{dat} the number of data points in the set which is added through reweighting.

To obtain the reweighted PDFs, all that has to be done is to compute the χ_k^2 of replica k to the inclusive jet data, using Eq. (4), for each of the $N_{\text{rep}} = 1000$ replicas of the NNPDF2.0 (DIS + DYP) parton set. For the inclusive jet data the total number of data points is $N_{\text{dat}} = 186$, and the covariance matrices are as given by the CDF and D0 Collaborations, the normalization uncertainty are incorporated using the t_0 -method, as discussed in Refs. [4,15]. The weight associated with each replica is then computed according to Eq. (9): specifically we evaluate

$$e_k \equiv \left(\frac{1}{2} N_{\text{dat}} - 1 \right) \log \chi_k^2 - \frac{1}{2} \chi_k^2 \quad (37)$$

hence if $\langle e_k \rangle \equiv \frac{1}{N_{\text{rep}}} \sum_{k=1}^{N_{\text{rep}}} e_k$, the weights are given by

$$w_k = \mathcal{N} \exp[e_k - \langle e_k \rangle], \quad \mathcal{N} = N_{\text{rep}} \left/ \sum_{k=1}^{N_{\text{rep}}} \exp[e_k - \langle e_k \rangle] \right. \quad (38)$$

The subtraction of $\langle e_k \rangle$ in the exponent is introduced to avoid numerical problems. We set to zero all weights for which $\exp[e_k - \langle e_k \rangle] < 10^{-12}$.

The χ_k^2 distributions for the jet data before and after reweighting, the $P(\alpha)$ estimator and the distribution of weights are shown in Fig. 1. We notice that before reweighting the distribution of χ^2 per data point is peaked close to one, but with a long tail extending to higher values of χ^2 . This has to be expected since the inclusive jet data are not included in the NNPDF2.0

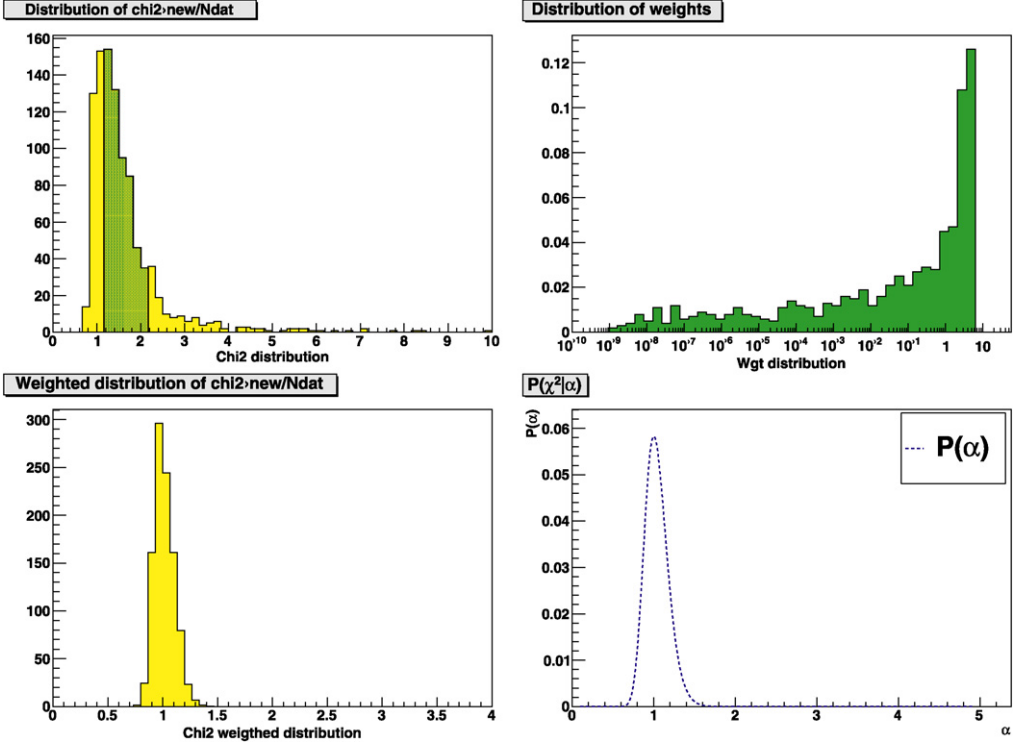


Fig. 1. Upper plots: distribution of χ_k^2/N_{dat} (the χ^2 per data point) and the weights w_k in the reweighting of the NNPDF20 (DIS + DYP) set using the inclusive jet data. Lower plots: distribution of the reweighted χ^2 distribution of the inclusive jet data, and the probability distribution $\mathcal{P}(\alpha)$ of the error rescaling parameter α .

(DIS + DYP) set. However 82% of the replicas have $0.5 < \chi_k^2 < 2$, confirming that the inclusive jet data are likely to be consistent with the other data in the fit and their inclusion in the fit will have only a moderate impact. Indeed a significant fraction of the weights are of order one, with however a long tail of small weights for replicas which will be effectively eliminated once the inclusive jet data are included.

To make these statements more quantitative, we can now evaluate the number of effective replicas, determined through the Shannon entropy according to Eq. (10): the effective number of replicas after reweighting using the jet data is $N_{\text{eff}} = 332$, i.e. around a third of the replicas are left.

To examine the consistency of the inclusive jet data with the DIS and Drell–Yan data used in NNPDF20 (DIS + DYP), we show in Fig. 1 the reweighted χ^2 distribution computed according to Eq. (11). Clearly the replicas which gave a poor fit to the jet data have now been removed, and the result is a distribution of χ^2 peaked at one, which shows that the jet data are consistent with the DIS and Drell–Yan data. This conclusion is reinforced by the probability distribution $\mathcal{P}(\alpha)$, plotted in Fig. 1: the most probable value for α is very close to one, showing that the overall size of the experimental errors of these data have been perfectly estimated by CDF and D0.

In order to determine quantitatively if indeed the refitted and reweighted PDF sets represent statistically identical distributions, we can compute the distances between central values and uncertainties of different PDF combinations, as defined in Ref. [4] with the required modifications

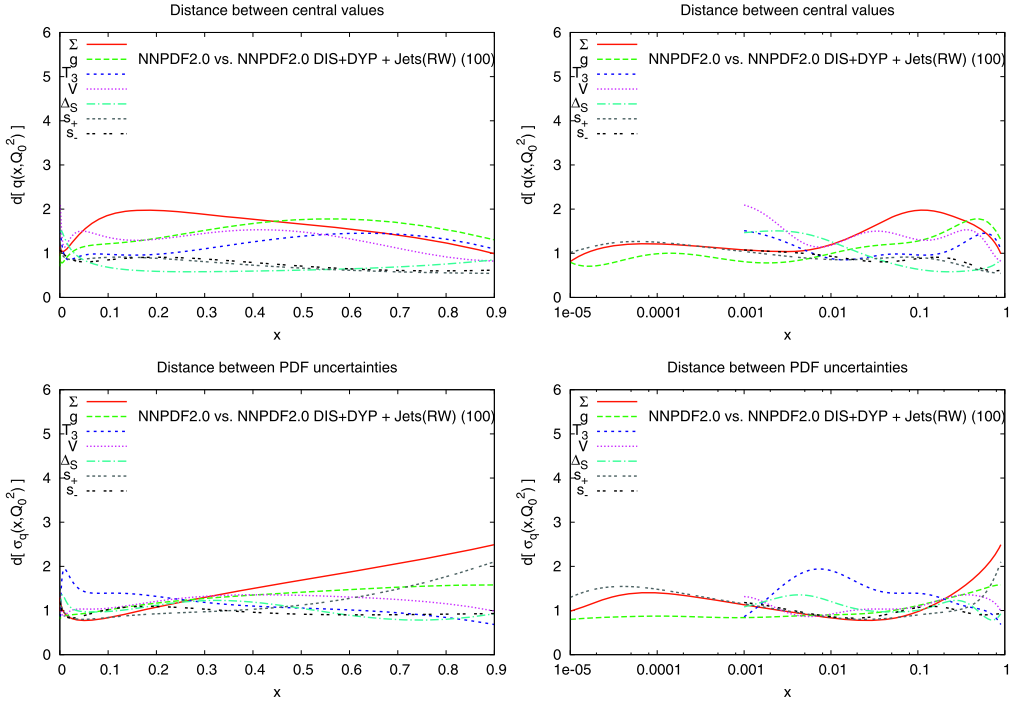


Fig. 2. Distances between PDFs (above) and uncertainties (below) for the NNPDF2.0 set and a set obtained adding the Tevatron inclusive jet production data to the NNPDF2.0 (DIS + DY) fit using the reweighting technique. The distances have been computed between sets of $N_{\text{rep}} = 100$ replicas.

to account for the individual weights of each replica.⁴ In Fig. 2 we plot the distances between PDFs' central values and uncertainties for the reweighted set and the (refitted) NNPDF2.0 set. Note that distances have been computed between sets of $N_{\text{rep}} = 100$ replicas. The distance is normalized such that distances $d \sim 1$ correspond to statistically identical distributions. We see that to a very good approximation the refitted and the reweighted sets are statistically equivalent, both for central values and uncertainties. The very largest distances, $d \sim 2$, corresponding a difference of about one seventh of a standard deviation of the measured quantity.

Given that as shown in Fig. 2 the refitted and reweighted sets are statistically equivalent, we know from [4] that inclusive jet data constrain only the large- x gluon, leaving virtually unchanged all other distributions. The reweighted gluon distribution and its uncertainty are shown in Fig. 3, compared with the original distribution, the NNPDF20 (DIS + DYP) fit, and with the full NNPDF2.0 fit. On the left-hand side we plot the gluon distribution with its uncertainty band and on the right-hand side the absolute value of the uncertainty. The reweighted and refitted distributions are indeed shown to be equivalent within errors. In particular the error of the medium and large- x gluon is sensibly reduced by the inclusion of the Tevatron inclusive jet data while, the other PDFs are essentially unchanged in both the refitted and the reweighted sets.

This statistical equivalence is an important check on the consistency of the NNPDF fitting methodology and the reweighting method presented here. In particular, it shows that an NNPDF

⁴ The expressions for the distances for reweighted PDF sets are collected in Appendix A.

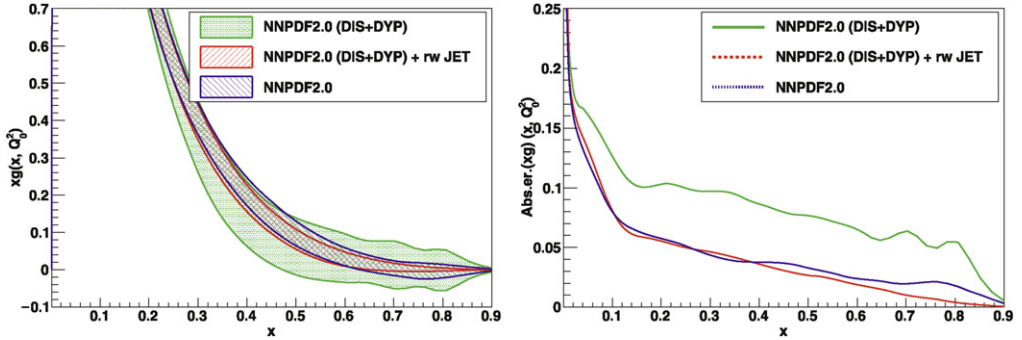


Fig. 3. The gluon distribution (left) and its uncertainty (right) of the NNPDF2.0 (DIS + DY) fit before and after reweighting with the inclusive jet data compared to the refitted gluon from NNPDF2.0 on a linear scale.

parton fit (at least in the case examined here) behaves in a way which is consistent with the laws of statistical inference: since reweighting is simply an application of probability theory, and since reweighting and refitting can be used interchangeably, the results obtained from the global fits indeed behave as probability distributions.

5. Application: The W lepton asymmetry

Now that we have explicitly verified that reweighting works, we can use it to assess the impact on PDF determination of data which were not included in the NNPDF2.0 fit. In this section we consider the Tevatron D0 W lepton charge asymmetry high luminosity data from Run II [8,9]. This data have attracted a lot of interest recently because of their potential inconsistency with other datasets which are traditionally included in the global fit like the deuterium DIS data [11, 12].

5.1. Motivation

In proton–antiproton scattering, W^\pm bosons are mainly produced by the annihilation of a u (d) quark in the proton with the \bar{d} (\bar{u}) in the antiproton. An asymmetry in the W^+ and W^- rapidity distributions is the result of a difference between the u and d distributions in the proton. Therefore, the information on the W charge asymmetry [19] provides a further constraint on the u and d PDFs. However, due to the unknown longitudinal momentum of the neutrino, the vector boson rapidity is difficult to determine directly. What is typically measured [7–10] is instead the lepton charge asymmetry. The vector boson rapidity may then be deduced in terms of the pseudo-rapidity of the charged lepton and its transverse energy E_T^l . Moreover, if the transverse energy E_T^l of the outgoing lepton is relatively small, the leading sea contribution $\bar{u} - \bar{d}$ is enhanced relative to the valence–valence contributions, so the lepton charge asymmetry also probes the separation into valence and sea quarks. For this reason in some experimental analysis [9,10], the lepton asymmetry is measured in different bins of E_T^l .

Historically, the Tevatron W lepton charge asymmetry data have been used in global fits together with the deuterium DIS data from BCDMS and NMC to constrain the ratio of d to u quarks at large- x . One advantage with respect DIS data is that theoretical uncertainties linked to the deuterium target, like nuclear effects, are not present for the lepton asymmetries, where only proton PDFs are involved. In Fig. 4 we show the d/u ratio computed using different recent

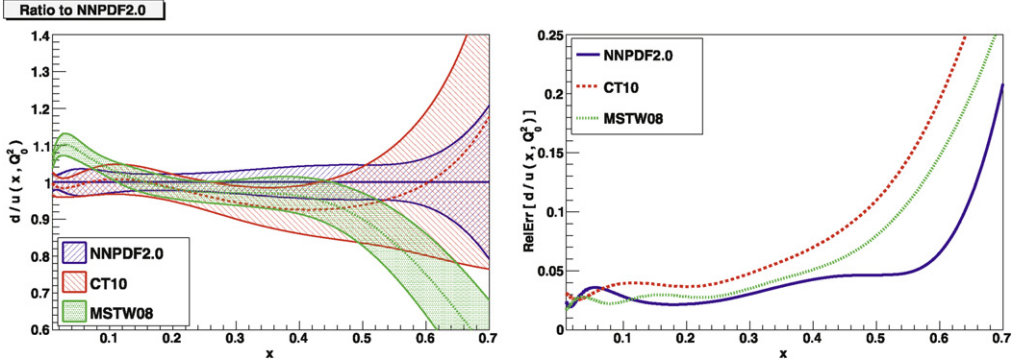


Fig. 4. The d/u ratio at large x computed at $Q_0^2 = 2 \text{ GeV}^2$ from the NNPDF2.0, MSTW08 and CT10 sets. We show the results for the ratio normalized to NNPDF2.0 (left plot) and the relative PDF uncertainties in each case (right plot). All uncertainties are 1σ .

PDF sets: NNPDF2.0, CT10 and MSTW08, together with the relative uncertainties. It is clear that PDF uncertainties are sizable for this combination at large- x , thus additional precision measurements of the W asymmetry are useful to reduce PDF errors in this region. We notice that in the kinematic region probed by the Tevatron measurements ($0.1 \lesssim x \lesssim 0.7$) the predictions from the three sets are in reasonable agreement within the respective uncertainties.

In the NNPDF2.0 analysis only the CDF W boson direct asymmetry data of Ref. [19] are included. This observable is implemented in the fitting code at next-to-leading order, without reverting to a K -factor approximation, using the FastKernel method described in [4]. The W lepton asymmetry measurements, on the other hand, were not included in the analysis due to the lack of a fast implementation. However, the recent development of the APPLGRID [20] interface is likely to facilitate the future inclusion of these data directly in our fits. Thanks to the reweighting technique presented earlier in this paper, we can now study the impact of the lepton asymmetry data consistently in NLO QCD.

Here we will consider the electron and muon asymmetry measurements performed by the D0 Collaboration at Run II of the Tevatron and published in Refs. [8,9]. The more recent D0 muon analysis of Ref. [10] has not been included since the data are still preliminary. The datasets included in our analysis are the same as those included in the dedicated CT10W analysis [11]. The lepton asymmetry measurements from CDF [7] are not considered here since the direct CDF W asymmetry data is already included in the NNPDF2.0 fit.

Let us discuss in more detail the lepton asymmetry data that we consider. In Ref. [8] a measurement of the muon charge asymmetry based on 0.3 fb^{-1} of data is presented. The asymmetry measurement is binned in ten bins in the muon pseudo-rapidity in the range $|\eta_\mu| < 2$, and cuts are imposed on the transverse energy and mass of the muon: $E_T^\mu > 20 \text{ GeV}$ and $M_T > 40 \text{ GeV}$. In Ref. [9] a similar measurement of the electron charge asymmetry is presented based on 0.75 fb^{-1} of data. The asymmetry is binned in twelve bins in the electron pseudo-rapidity in the range $|\eta_e| < 3.2$, and cuts are imposed on the missing energy and transverse mass: $\cancel{E} > 25 \text{ GeV}$ and $M_T > 50 \text{ GeV}$. Three sets of measurements are then given, which have different cuts in the transverse energy of the electron: an ‘inclusive’ bin which has $E_T^e > 25 \text{ GeV}$ (which we refer to here as bin A), and two less inclusive bins with more restrictive cuts on the transverse energy, $25 \text{ GeV} < E_T^e < 35 \text{ GeV}$ (bin B) and $E_T^e > 35 \text{ GeV}$ (bin C). Note that bins B and C together cover the same kinematic range as the more inclusive bin A.




Fig. 5. Predictions for the D0 W lepton charge asymmetry obtained with the DYNNLO code at next-to-leading order, using the NNPDF2.0 [4], CT10 [11] and MSTW08 [21] parton sets. We show results for the muon charge asymmetry (top left), and the electron charge asymmetry in the inclusive bin, $E_T^e > 25$ GeV, bin A (top right), and then in less inclusive bins, $25 \text{ GeV} < E_T^e < 35$ GeV, bin B (bottom left), and $E_T^e > 35$ GeV, bin C (bottom right).

To analyse these data using the reweighting technique we use the DYNNLO code [22] to compute the theoretical predictions for the lepton asymmetries at NLO, using NNPDF2.0 as input parton densities. This code is a parton level Monte Carlo program designed to compute exclusive hadronic processes up to NNLO, and it enables the user to implement the same cuts used in the experimental analyses.

Before considering reweighting, let us first compare in Fig. 5 the predictions obtained with DYNNLO and different PDF sets at NLO, for the various D0 datasets. It is perhaps surprising that, even though none of these data are included into the NNPDF2.0 fit, the prediction obtained from the NNPDF2.0 is in general closer to the experimental data than the predictions obtained with the other parton sets: the reason can be traced back to the somewhat larger d/u ratio in the range $0.2 \lesssim x \lesssim 0.6$ (Fig. 4) for the NNPDF2.0 set. The exception is bin B, for which MSTW08 provides the best description.

The quality of the comparison of various PDF sets with the asymmetry data can be quantified by evaluating the χ^2 to each data set. For all the Tevatron Run II D0 lepton asymmetry only the statistical and uncorrelated systematic errors are quoted. The covariance matrix is therefore diagonal and its elements are given by the sum in quadrature of the statistical and the uncorrelated systematic errors. There is no normalization uncertainty since the asymmetry is a ratio of cross-sections.

Table 1

The D0 W lepton charge asymmetry datasets that are included in the present analysis, together with the χ^2 per data point obtained from the NLO predictions of various PDFs sets. The electron data of Ref. [9] is divided into three bins that we denote by bin A, bin B and bin C.

Set	N_{dat}	NNPDF2.0	MSTW08	CT10
D0 μ ($E_T^\mu > 20$ GeV) [8]	10	0.62	1.51	0.70
D0 electron, $E_T^e > 25$ GeV (bin A) [9]	12	2.12	9.20	4.07
D0 electron, $25 \text{ GeV} < E_T^e < 35$ GeV (bin B) [9]	12	4.75	1.66	9.48
D0 electron, $E_T^e > 35$ GeV (bin C) [9]	12	5.06	13.4	11.7

The value of the χ^2 per data point and the number of data points for each set considered in the present analysis are shown in Table 1. The results confirm the studies performed in Ref. [23]. In particular, the less inclusive data (bins B and C) are rather poorly described by all the current PDF fits, with the exception of bin B which MSTW08 describes reasonably well (though at the cost of a very bad fit to bins A and C). Note however that Ref. [11] uses the RESBOS program [24] to compute the predictions for the W lepton asymmetry. RESBOS computes on top of the NLO higher order corrections from p_T resummation. The differences between NLO and RESBOS are maximal in the kinematics of the electron bin B data. This differences might explain, at least in part, the values of the χ^2 for CT10 obtained in Table 1 compared to those given in Ref. [11].

We now consider the effect of including the Run II D0 muon and electron asymmetry data in the NNPDF2.0 analysis using reweighting. We will consider each dataset in turn, concentrating first on the inclusive sets (muon and electron bin A), and turning later to the less inclusive data sets (bins B and C). For each case we will proceed as follows: first we provide the distribution of χ_k^2 before and after reweighting, the probability distribution $P(\alpha_s)$ and the distribution of weights. We then compare the reweighted PDFs to experimental data. Finally we compute the distances between the original and the reweighted sets, and compare the corresponding PDFs where they differ substantially from the original ones.

Unless otherwise stated, PDFs, their uncertainties will be plotted at the scale $Q^2 = Q_0^2 = 2 \text{ GeV}^2$. The $N_{\text{rep}} = 1000$ NNPDF2.0 set is used throughout, with the exception of the computation of distances, where we instead use sets of 100 replicas.

5.2. Inclusive data

Let us first consider the inclusion of the D0 muon charge asymmetry data [8]. The distribution of the χ_k^2 and corresponding weights w_k for these data is shown in the upper plots of Fig. 6. Since the χ^2 -distribution is peaked close to one, the weights are also mostly of order unity. The reweighted χ^2 , and probability density for the rescaling parameter α are shown in the lower plots: they peak rather below one, suggesting that the errors on these data are actually likely to be overestimated by D0 by around 20%. After reweighting the χ^2 per data point drops from 0.62 to 0.51, and the number of effective replicas is $N_{\text{eff}} = 795$.

On the left in Fig. 7 we show the muon charge asymmetry before and after the reweighting. Indeed the predictions get closer to the data, once the PDFs are reweighted. We have also examined the effect on the shape of the PDFs, but the effects are negligible apart from a slight reduction in the uncertainty of the total valence distribution, shown in Fig. 8. This is confirmed by the distance analysis, Fig. 9, that shows that central values and uncertainties for all PDFs are essentially unchanged, with the exception of the total valence PDF where the inclusion of muon data has a moderate effect.

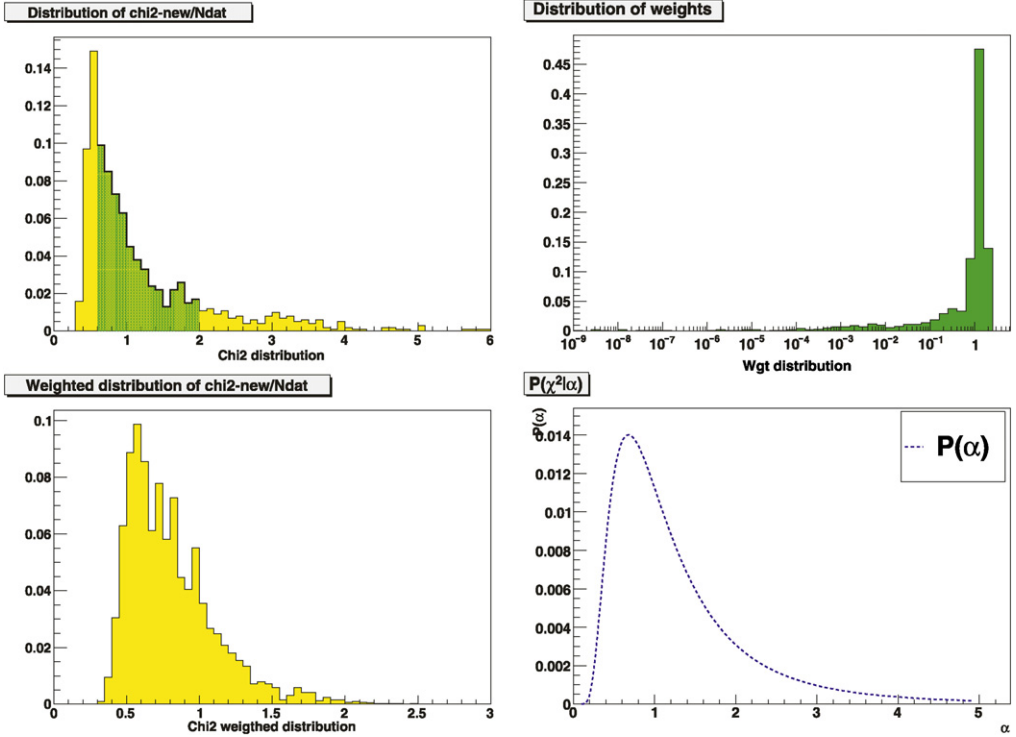


Fig. 6. Distribution of the χ_k^2 and the weights w_k , the reweighted χ^2 -distribution and the probability distribution $\mathcal{P}(\alpha)$ in the reweighting of the NNPDF2.0 PDF set using the D0 muon asymmetry data [8].




Fig. 7. Left: W muon charge asymmetry computed for the NNPDF2.0 PDFs before and after the reweighting of these data into the parton analysis. Right: W electron charge asymmetry (inclusive bin) computed with the NNPDF2.0 PDFs before and after the reweighting of these data in the parton analysis.

To study the compatibility of these data with the data included in the NNPDF2.0 analysis, in Table 2 we show the χ^2 of each of the datasets included in the NNPDF2.0 analysis evaluated with the original NNPDF2.0 PDFs and then with these PDFs reweighted by the inclusion of the D0 muon asymmetry data. If anything, there is a slight improvement in the description of




Fig. 8. Total valence PDF for the NNPDF2.0 and NNPDF2.0 + D0 muon data PDF sets.

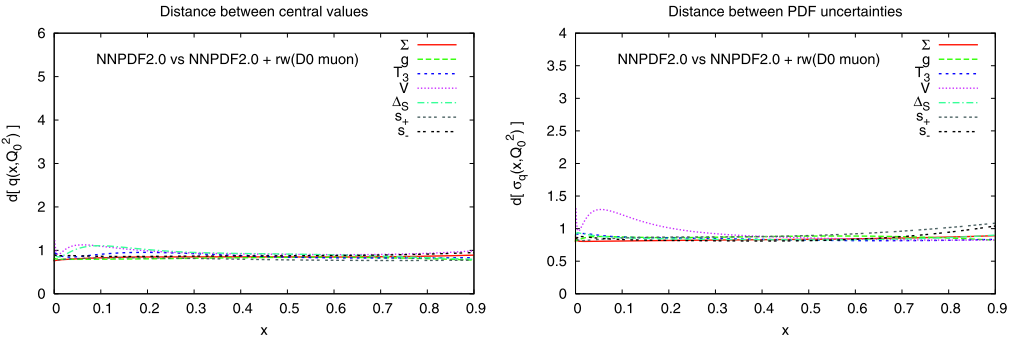


Fig. 9. Distances between NNPDF2.0 and NNPDF2.0 + D0 W lepton asymmetry measurements from the muon dataset. The NNPDF2.0 set with $N_{\text{rep}} = 100$ has been used in the computation of the distances.

most of the datasets. To summarize, the D0 muon asymmetry data [8] are perfectly consistent with NNPDF2.0, but are not sufficiently precise to add much information to the PDFs. It will be interesting to assess the impact of the higher statistics D0 Run II muon data set [10] once the analysis is completed.

Next we consider the inclusive D0 electron data (bin A) with $E_T^e > 25$ GeV. The results are shown in Fig. 10. Once included in the fit through reweighting the χ^2 for this set drops from 2.12 to 1.55. While the distribution of the unweighted χ_k^2 is peaked above two and has a long tail to higher values, after reweighting the peak is shifted much closer to one. This is achieved through a substantial reduction in the effective number of replicas: after reweighting $N_{\text{eff}} = 262$. However, while before reweighting only 16% of replicas lie in the region $\frac{1}{2} < \chi^2 < 2$, after reweighting this figure rises to 78%. This behavior is confirmed by the plot of $\mathcal{P}(\alpha)$: the data indicate that the most probable value of α is around 1.7, indicating that experimental errors on these data are a little underestimated (by around 30%). Taken together, this shows that these data might have a significant effect on constraining the PDFs, while still being broadly consistent with all the other data included in the fit.

The improvement in the description of the electron asymmetry after reweighting in Fig. 7, while the fit to the other datasets included in the NNPDF2.0 fit shows no significant deterioration: if any change has to be noticed, is a slight improvement in particular in the fit to the CDF W asymmetry data (see Table 2).

Table 2

χ^2 per data point of all the experiments included in the NNPDF2.0 fit evaluated before and after reweighting with the various lepton asymmetry data sets. Note that here we use the t_0 covariance matrix in the evaluation of the χ^2 : the numbers are thus slightly different from those shown in Ref. [4]. The cases in which the χ^2 varies significantly as compared to the reference are highlighted in boldface.

Set	$\chi^2_{2.0}$	$\chi^2_{2.0+\mu}$	$\chi^2_{2.0+\text{binA}}$	$\chi^2_{2.0+\text{binA}+\mu}$	$\chi^2_{2.0+\text{binB}}$	$\chi^2_{2.0+\text{binC}}$
NMC-pd	0.99	0.98	0.98	0.98	0.97	1.13
NMC	1.72	1.72	1.69	1.70	1.72	1.72
SLACp	1.55	1.55	1.53	1.54	1.50	1.63
SLACd	1.12	1.12	1.07	1.09	1.05	1.24
BCDMSp	1.35	1.35	1.33	1.34	1.41	1.35
BCDMSd	1.16	1.16	1.16	1.16	1.24	1.14
HERA1-NCep	1.35	1.35	1.34	1.34	1.33	1.35
HERA1-NCem	0.86	0.86	0.86	0.86	0.86	0.86
HERA1-CCep	0.96	0.96	0.94	0.94	1.02	0.92
HERA1-CCem	0.56	0.56	0.56	0.56	0.56	0.57
CHORUSnu	1.08	1.08	1.08	1.08	1.11	1.10
CHORUSnb	0.86	0.86	0.86	0.86	0.87	0.90
FLH108	1.50	1.50	1.50	1.50	1.47	1.50
NTVnuDMN	0.69	0.66	0.67	0.65	0.82	0.60
NTVnbDMN	0.70	0.70	0.69	0.69	0.72	0.81
Z06NC	1.24	1.24	1.24	1.24	1.23	1.26
Z06CC	1.19	1.19	1.19	1.19	1.15	1.21
DYE605	0.86	0.86	0.84	0.85	0.87	0.85
DYE886p	1.31	1.32	1.29	1.30	1.28	1.36
DYE886r	0.83	0.79	0.67	0.71	1.08	0.72
CDFWASY	1.88	1.88	1.78	1.82	2.05	1.60
CDFZRAP	1.74	1.77	1.75	1.77	1.37	1.97
D0ZRAP	0.59	0.59	0.59	0.59	0.60	0.61
CDFR2KT	1.02	1.02	0.95	0.97	1.21	0.93
DOR2CON	0.86	0.86	0.84	0.84	0.91	0.84
TOTAL	1.14	1.14	1.13	1.13	1.16	1.16

In Fig. 12 we plot the distances between the prior set NNPDF2.0 and the reweighted set: it is clear that the most significant effect is on the uncertainty in the valence PDF. Indeed, in Fig. 11 we show the error reduction that comes from the inclusion of the inclusive D0 electron charge asymmetry data on the valence PDF. While the central value remains essentially unchanged, the uncertainty is significantly reduced. Small improvements in the precision of the singlet and triplet quark distributions can also be observed, while other PDFs combination remain unchanged.

Having found that both the inclusive muon and electron (bin A) D0 asymmetry data are each consistent with the datasets used in NNPDF2.0, it is interesting to ask whether they are also consistent with each other. This is not obvious a priori: it is in principle possible for each dataset to prefer a different subset of the NNPDF2.0 replicas.

To examine this question we performed a reweighting analysis with the combined dataset: the χ^2_k used to determine the weights are then the sum of those from the D0 muon asymmetry data and the D0 electron asymmetry data, i.e. $N_{\text{dat}} = 22$ data points. The number of effective replicas is then reduced to $N_{\text{eff}} = 356$, actually a number larger than the case where electron data alone were considered: the muon data soften the impact of these data. The combined χ^2_k and α distributions (see Fig. 13) are now better behaved: while before the reweighting only 49% of replicas have a $\frac{1}{2} < \chi^2 < 2$, after reweighting this now rises to 99%. The peak of the α distribution is

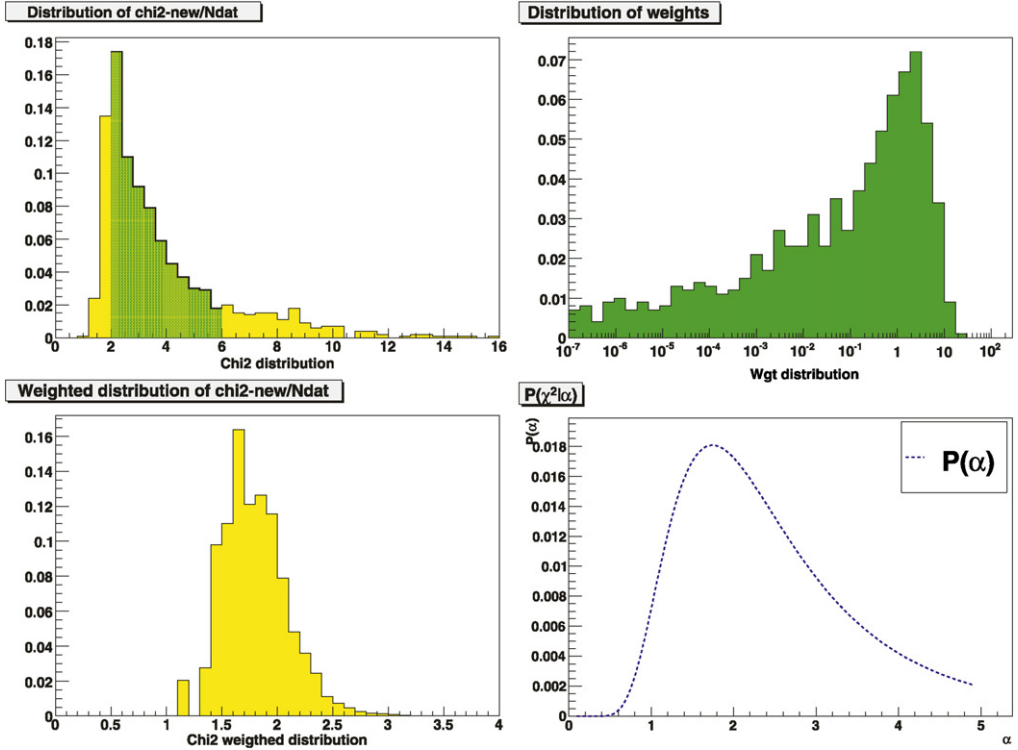


Fig. 10. Distribution of the χ_k^2 and the weights w_k , the reweighted χ^2 -distribution and the probability distribution $\mathcal{P}(\alpha)$ in the reweighting of the NNPDF2.0 PDF set using the D0 electron asymmetry data (bin A) [9].

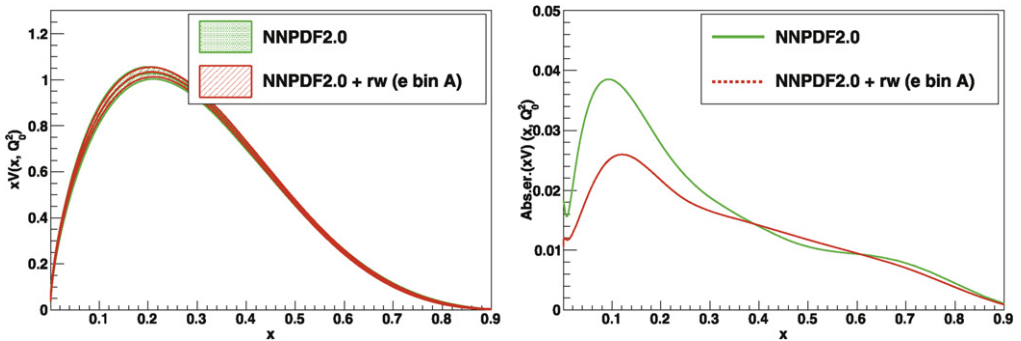


Fig. 11. Total valence PDF for NNPDF2.0 and NNPDF2.0 + D0 electron data (bin A).

now quite close to one: the overestimated uncertainties of the muon data complement the underestimated uncertainties of the electron data. The quality of the fit to the other datasets included in the NNPDF2.0 fit shows no significant deterioration, and again there is a slight improvement, in the fit to the CDF W asymmetry data (see Table 2).

In Fig. 14 we show the effect of the addition of the D0 muon and the D0 electron inclusive data on the valence distribution. The precision of the valence distribution is significantly im-

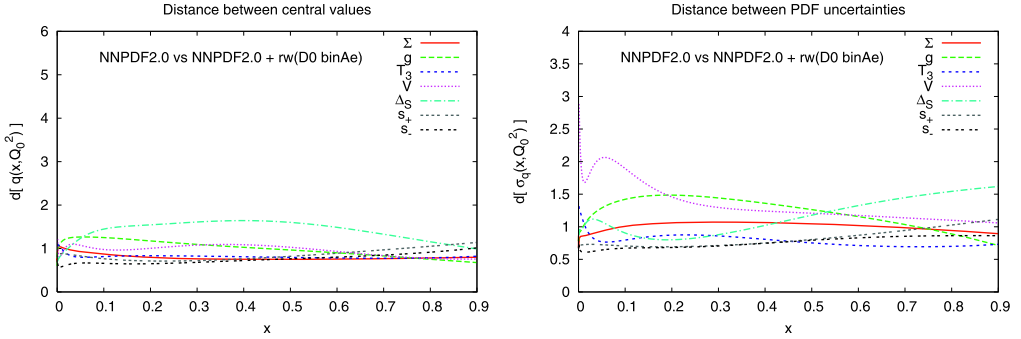


Fig. 12. Distances between NNPFD2.0 and NNPFD2.0 + D0 W lepton charge asymmetry measurements from the electron bin A dataset. The NNPFD2.0 set with $N_{rep} = 100$ has been used in the computation of the distances.

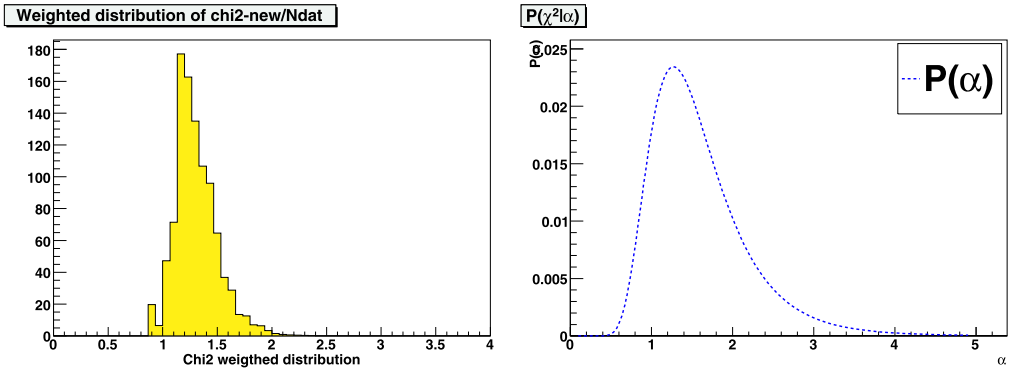


Fig. 13. The reweighted χ^2 -distribution and the probability distribution $\mathcal{P}(\alpha)$ in the reweighting of the NNPFD2.0 PDF set using the combined D0 muon asymmetry data and electron asymmetry data (bin A).

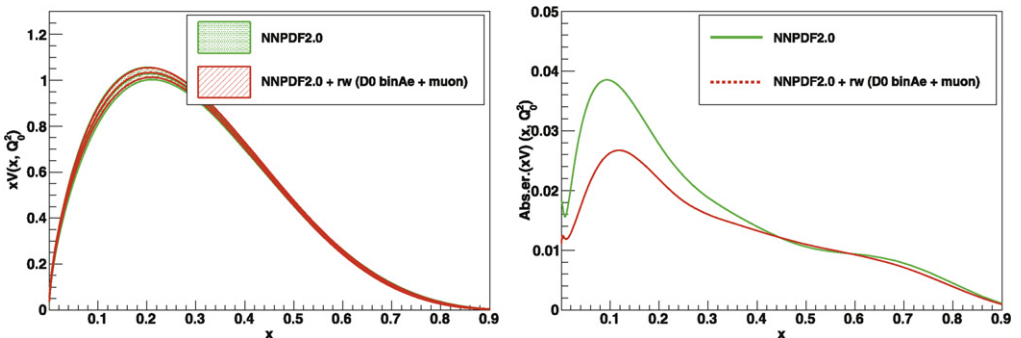


Fig. 14. Total valence PDF for the NNPFD2.0 and NNPFD2.0 + D0 muon + D0 electron data (bin A) sets.

proved, though without shifting its central value significantly. This implies that the NNPFD2.0 set is quite consistent with the inclusive data, so that their addition entails only PDF uncertainty reduction without affecting central values. It follows that the d/u ratio extracted from the DIS

Table 3

A summary of the results of reweighting with the D0 W lepton asymmetry data: the fraction $N_{\text{eff}}/1000$ of replicas left after reweighting, the most probable value α_{opt} of the error rescaling parameter α , the χ^2 per data point to the D0 W lepton data evaluated before and after the reweighting, and the total χ^2 per data point to all the other data in the NNPDF2.0 fit.

Set	$N_{\text{eff}}/1000$	α_{opt}	χ^2	χ_{rw}^2	$\chi_{\text{tot-rw}}^2$
D0 μ ($E_T > 20$ GeV)	0.795	0.7	0.62	0.51	1.14
D0 ebin A ($E_T > 25$ GeV)	0.262	1.7	2.12	1.55	1.13
D0 μ + ebin A	0.356	1.3	1.44	1.11	1.13
D0 ebin B ($25 \text{ GeV} < E_T < 35 \text{ GeV}$)	0.061	1.3	4.75	1.12	1.16
D0 ebin C ($E_T > 35$ GeV)	0.068	2.7	5.06	2.51	1.16

deuterium data and the CDF direct W charge asymmetry data will be consistent with the information included in the D0 inclusive muon and electron data.

The main statistical estimators for the lepton charge asymmetry data sets are summarized in Table 3. The two inclusive sets have a significant impact on PDFs, and are reasonably consistent with themselves (though the experimental uncertainties on the inclusive D0 electron charge asymmetry data, bin A, may be a little underestimated), with the other data used in the NNPDF2.0 fit, in particular the CDF W charge asymmetry data, and with each other.

The χ^2 values for the total dataset and for the individual experiments in the NNPDF2.0 analysis are shown in Table 2. The sets that differ sizably from the reference results have been highlighted in boldface in the different cases. As far as the inclusive muon and electron datasets are concerned we notice that both are consistent with the NNPDF2.0 datasets, and their inclusion improves the fit to the W asymmetry data. Furthermore they are both consistent with each other. These conclusions do not support the conclusions of the MSTW08 analysis [12], which finds that inclusion of the D0 electron inclusive bin in the global fit, without significant deterioration in the fit to the other datasets, requires sizable nuclear corrections to deuterium data.

5.3. More exclusive data

We now turn to the less inclusive D0 electron charge asymmetry data (bin B and bin C), where the transverse energy of the electron is restricted to the range $25 \text{ GeV} < E_T^e < 35 \text{ GeV}$ and $E_T^e > 35 \text{ GeV}$ respectively.

We first consider each bin separately, and we turn then to their combination. Considering first the lower E_T^e bin (bin B), the number of effective replicas is now reduced to $N_{\text{eff}} = 61$, indicating that, as expected, these data are more constraining than those of the inclusive bin. This is so because the data binned in E_T^e probe a more localized region in x of the PDFs as compared to the inclusive data. The χ^2 for this set drops from 4.75 to 1.12 after the data is included. From the plots in Fig. 15 we see that indeed there is now a significant fraction of very small weights, because many of the replicas fit the new data rather badly. However after reweighting the χ^2 distribution improves very significantly: while before reweighting only 4.8% of replicas were in the range $\frac{1}{2} < \chi^2 < 2$, after reweighting this increases to 86.5%. The rescaling plot peaks at around $\alpha \sim 1.3$ indicating that the errors on the data are reasonably well estimated.

The improvement in the fit to the lowest E_T^e bin charge asymmetry data is manifest on the left of Fig. 18. However the fit to some of the other datasets in the NNPDF2.0 fit, in particular BCDMSp and BCDMSd, becomes significantly worse (see Table 2). The fact that there is as much tension here with the proton data as with the deuteron data suggests that it is unlikely

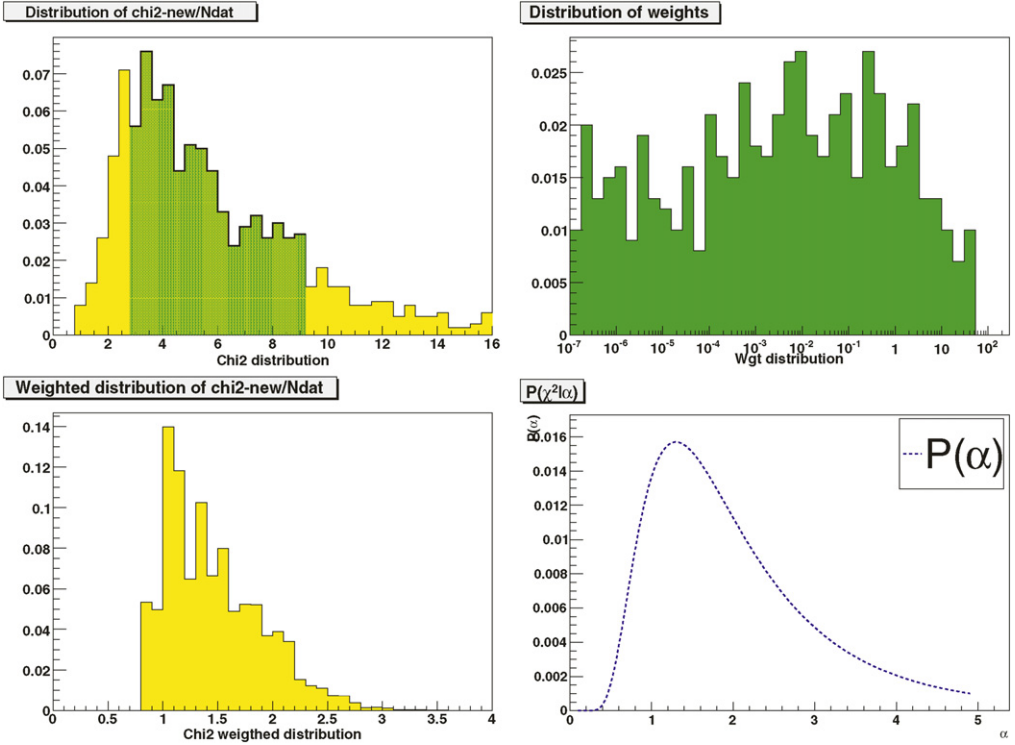


Fig. 15. Distribution of the χ_k^2 and the weights w_k , the reweighted χ^2 -distribution and the probability distribution $\mathcal{P}(\alpha)$ in the reweighting of the NNPDF2.0 PDF set using the D0 electron asymmetry data bin B [9].

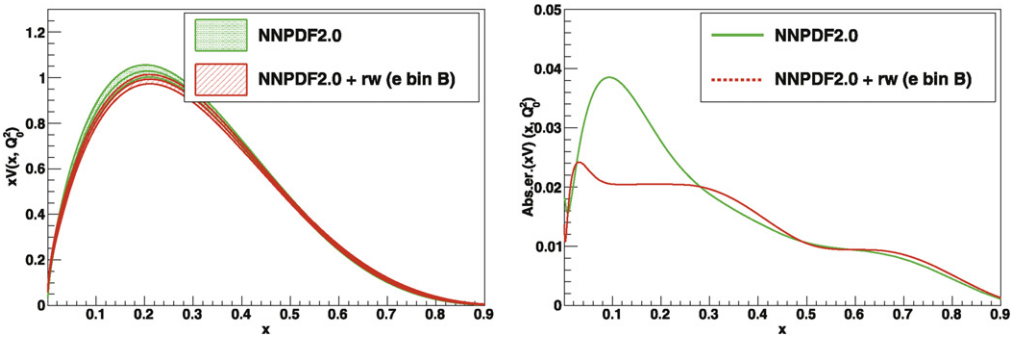


Fig. 16. Total valence PDF for NNPDF2.0 and NNPDF2.0 + D0 electron data (bin B).

that nuclear corrections to the deuteron target can help (contrary to the claim in Ref. [12]). The overall χ^2 per degree of freedom rises from 1.14 to 1.16: this is rather significant, given that we are only adding 12 new data points to the 3415 used in NNPDF2.0. The decrease in the fit quality is driven by the large weight that BCDMS carry in the global fit. It should be further noted that the fit to the inclusive jet data and CDF W asymmetry also worsens.

These problems are also apparent when we look at the effect on individual PDFs: in particular while the valence distribution, Fig. 16, is now better determined in some ranges of x , elsewhere

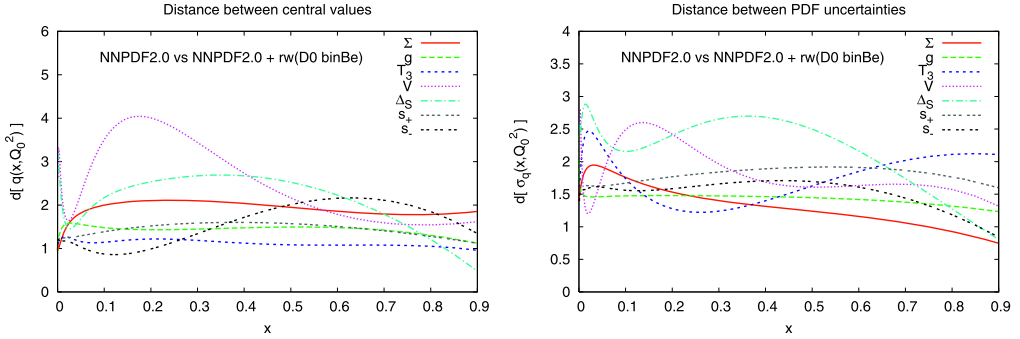


Fig. 17. Distances between NNPDF2.0 and NNPDF2.0 + D0 W lepton asymmetry measurements from the electron bin B dataset. The NNPDF2.0 set with $N_{\text{rep}} = 100$ has been used in the computation of the distances.




Fig. 18. W electron asymmetry computed on the NNPDF2.0 set before and after the reweighting of the D0 W electron asymmetry: bin B (on the left) and bin C (on the right).

the uncertainty increases. While this may in part be due to the rather limited statistics of the reweighted distribution, it is probably also a sign of some inconsistency with the other data used in the NNPDF2.0 fit. The statistical distances, plotted in Fig. 17, are also sizable for some PDFs, especially the valence distribution.

We finally consider the remaining D0 electron asymmetry dataset at highest E_T (bin C): the results are displayed in Fig. 19. While the impact of these data is similar to that of the lowest E_T bin (bin B), with the effective number of replicas dropping to 68, the quality of the fit to the unweighted replicas is so poor (there are no replicas with a χ^2 below 2) that even after reweighting the quality of the fit is still not very good, the average χ^2 per data point dropping from 5.06 to 2.51.

The rescaling plot shows a preferred value of $\alpha \sim 2.7$, suggesting that the experimental errors in these data are seriously underestimated. This might be caused by some underestimated systematic uncertainties in the separation of the data into bins of different E_T^e . The poor quality of the fit, even after reweighting, is again apparent in Fig. 18: it is clear that some of the bins simply cannot be fitted with a reasonably smooth distribution. There is also tension between these data and some of the other datasets included in NNPDF2.0 (see Table 2): in particular while BCDMS is now fine, the fit to the NMC-pd ratio is spoiled.

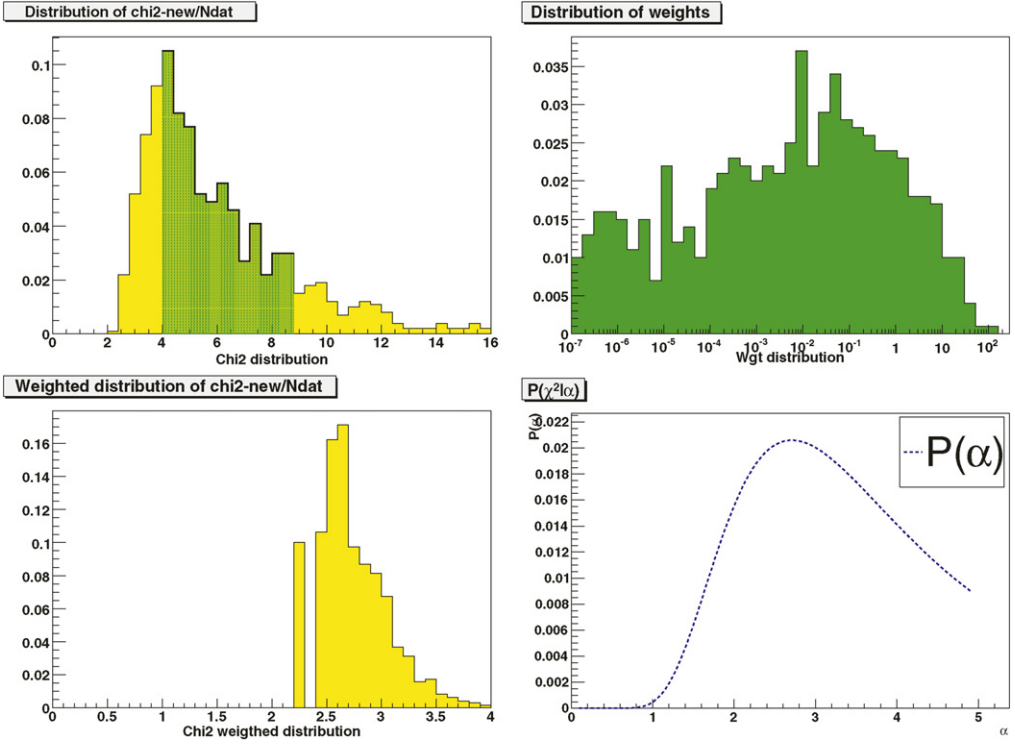


Fig. 19. Distribution of the χ_k^2 and the weights w_k , the reweighted χ^2 -distribution and the probability distribution $\mathcal{P}(\alpha)$ in the reweighting of the NNPDF2.0 PDF set using the D0 electron asymmetry data bin C [9].

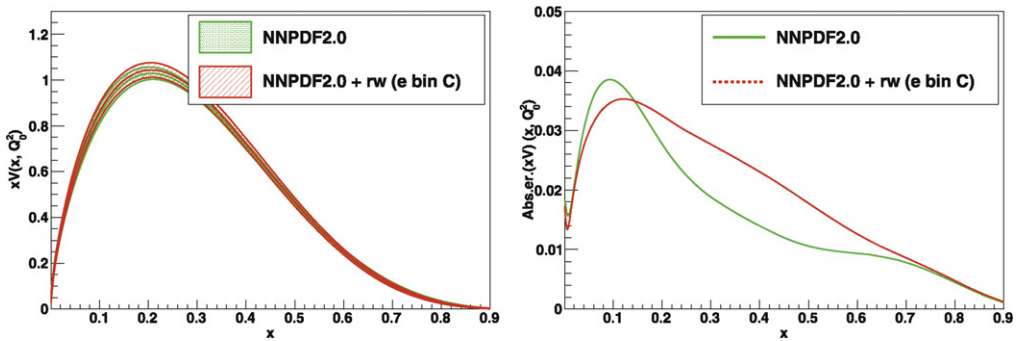


Fig. 20. Total valence for NNPDF2.0 and NNPDF2.0 + D0 electron data (bin C).

When we examine the effect of these data on the PDFs, we see (Fig. 20) that rather than making the PDFs more precise, in many regions of x the uncertainty increases substantially. The enlarging of the uncertainty is of course what one would expect when inconsistent data are combined, and it was previously seen to occur in NNPDF parton fits (see e.g. Section 3.4.1 of Ref. [25]). Here, it is shown to occur as a consequence of standard statistical inference.

We also attempted a combined fit of the D0 electron asymmetry data bins B and C, but the constraint imposed on PDFs by including these data together is so severe that the number of

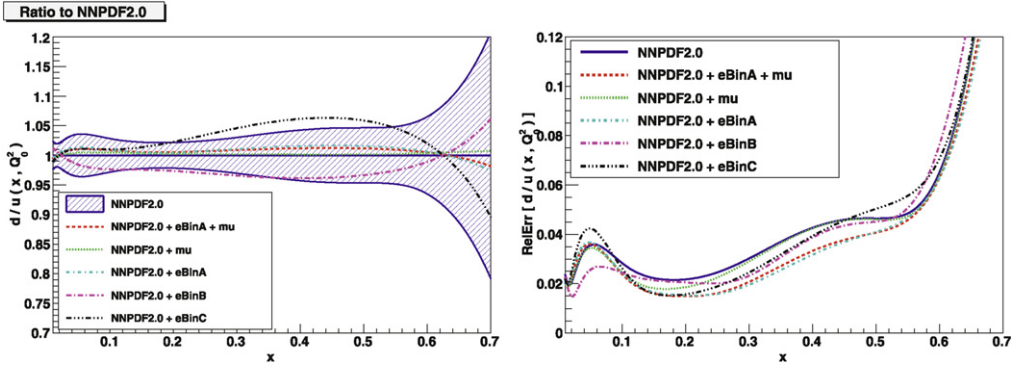


Fig. 21. The d/u ratio at large x computed at $Q_0^2 = 2 \text{ GeV}^2$ from the original NNPDF2.0 sets and the various sets obtained through reweighting of NNPDF2.0. We show the results for the ratio normalized to NNPDF2.0 (left plot) and the relative PDF uncertainties in each case (right plot). All uncertainties are 1σ .

effective replicas is reduced to one. This shows that not only are these data each inconsistent with other data included in the global fit, but they are also inconsistent with each other.

The main statistical estimators for the exclusive electron charge asymmetry data sets are summarized in Table 3. In contrast to what we observe for the inclusive sets, these data sets, while having an even greater effect on the PDFs, appear to be internally inconsistent (bin C), inconsistent with other data used in NNPDF2.0, particularly BCDMS proton and deuteron data (bin B), and also inconsistent with each other.

The results in the present study cannot be directly compared to the ones obtained in the CT10 analysis, because there the three electron E_T^e bins are added simultaneously to the fit. On top of the double counting problem, this is problematic because internal tensions of experimental origin between the different bins might be mistaken for a physical effect, such as nuclear corrections. Indeed, we have shown that the more exclusive data sets are not only inconsistent with other sets in the global analysis but also inconsistent among themselves, so that it is probably not a good idea to include both in the fit simultaneously.

5.4. Implications for the d/u ratio, and LHC benchmarks

Up to now we have considered the impact of the D0 data on different PDF combinations, noticing that the most relevant effect was on the total valence distribution. To conclude our analysis we assess the impact of W lepton charge asymmetry data on the d/u ratio. In Fig. 21 we display the d/u ratio at large x computed at $Q_0^2 = 2 \text{ GeV}^2$ from the original NNPDF2.0 set and the four sets obtained through reweighting of NNPDF2.0 with the D0 lepton asymmetry data. The effect of the inclusive datasets (muon and electron bin A) is rather small, even when they are combined together. The less inclusive sets (bin B and bin C) have a rather larger effect, but pull in opposite directions. Even so, the effect is only of the same order as the PDF uncertainty.

In Fig. 22 we compare the d/u ratio obtained with NNPDF2.0 and with NNPDF2.0 reweighted by the maximally consistent combination of the D0 data (muons and electrons bin A) with the CT10 and CT10W results, normalized to NNPDF2.0. It can be seen that the combination of D0 muon and electron bin A data leads to a substantial error reduction of $\sim 25\%$ in the d/u ratio in the $0.1 \lesssim x \lesssim 0.5$ region, with almost no change in the central value. Note also that the d/u ratio obtained from the NNPDF2.0 + D0 (ebin A + μ) set is rather more precise than that

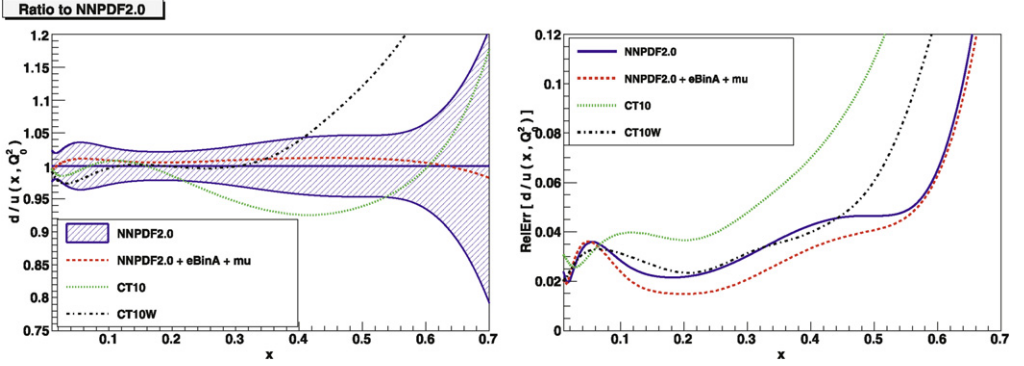


Fig. 22. Lower plots: The same d/u ratio from the original NNPDF2.0 sets, the NNPDF2.0 set reweighted by the maximally consistent combination of the D0 lepton asymmetry data (the muon data plus the inclusive electron data) and the CT10 and CT10W sets.

Table 4

Cross sections for different Standard Candle processes at the LHC (7 TeV) computed using NNPDF2.0 reweighted PDFs including D0 W lepton asymmetry data. The Higgs cross-section is computed for $m_h = 120$ GeV.

	$\sigma(Z \rightarrow ll^-)$	$\sigma(W^- \rightarrow l\nu)$	$\sigma(W^+ \rightarrow l\nu)$	$\sigma(h^0)$	$\sigma_{t\bar{t}}$
2.0	911 ± 16 pb	3.98 ± 0.08 nb	5.80 ± 0.12 nb	11.59 ± 0.18 pb	169 ± 6 pb
2.0 + D0 (μ)	911 ± 16 pb	3.98 ± 0.08 nb	5.80 ± 0.12 nb	11.58 ± 0.19 pb	169 ± 6 pb
2.0 + D0 (bin A)	914 ± 15 pb	4.00 ± 0.07 nb	5.81 ± 0.11 nb	11.58 ± 0.22 pb	168 ± 6 pb
2.0 + D0 (μ + bin A)	913 ± 15 pb	4.00 ± 0.07 nb	5.81 ± 0.11 nb	11.58 ± 0.20 pb	168 ± 5 pb
2.0 + D0 (bin B)	904 ± 15 pb	3.92 ± 0.07 nb	5.78 ± 0.11 nb	11.66 ± 0.15 pb	172 ± 5 pb
2.0 + D0 (bin C)	913 ± 22 pb	4.01 ± 0.10 nb	5.78 ± 0.17 nb	11.52 ± 0.28 pb	168 ± 8 pb

from CT10W, despite the fact that they include all the D0 lepton datasets, with larger weights than the other datasets in the global analysis.

Finally, it is interesting to ask to what extent the inclusion of the W lepton charge asymmetry data through reweighting affects the determination of some of the LHC standard cross-sections. Results for vector boson production, Higgs and $t\bar{t}$ at $\sqrt{s} = 7$ TeV are collected in Table 4. They have been computed using MCFM [26–28] to determine the cross-section for each replica, and then the weighted average of the results evaluated using Eq. (2). The uncertainties in each case are purely PDF uncertainties obtained from a reweighted evaluation of the variance of the cross-section. Clearly all these cross-sections are by and large insensitive to the addition of the D0 lepton charge asymmetry data, even those data (bins B and C) which show inconsistencies with the global dataset and thus have the largest (though least reliable) effect. This is to be expected since the LHC observables we have considered are not directly sensitive to large- x quarks, for which the impact of the D0 data is the largest.

6. Conclusions

In this paper we have developed a method for determining the effect of new data on PDFs without the need for a global refitting. The method relies on the existence of an ensemble of PDFs, distributed according to the uncertainties in a global set of older data, and thus representing the prior probability distribution of the PDFs. Such ensembles are provided by the NNPDF

Collaboration. The effect of new data is then accounted for by reweighting the PDF replicas in the ensemble according to their relative probabilities given the new dataset. These probabilities are determined simply and easily by computing the χ^2 of the new data to the prediction obtained using a given replica.

We have provided a careful derivation of our formula used to determine the weights. This is important because our result differs from that obtained in a previous attempt to use a reweighting method [5]. The derivation is subtle because it is necessary to deal with multi-dimensional probability densities, where unless one is careful one can fall into inconsistencies due to the Borel–Kolmogorov paradox [16].

The main advantage of the new method is clear: computing the weights is no more difficult or computer intensive than the usual procedure of preparing a plot comparing the new dataset with predictions from given PDFs. However the information provided is much more substantial – one can assess quantitatively the impact of the new data on the PDFs, whether the new data are consistent with all the older data encoded within the PDF ensemble and the theoretical assumptions on which it was based, and then whether the new data have any effect on other observables of interest such as benchmark cross-sections. Only when the impact of the new data is very large does a full refitting of the PDF ensemble become necessary, due to the loss of efficiency in the reweighted ensemble.

We thus envisage our method being useful to experimentalists in all sorts of situations: testing the reliability of preliminary datasets and their uncertainties, assessing the credibility of possible indications of new physics, or in optimizing the design of new experiments using pseudodata.

We have shown explicitly that the method works by considering the addition of Tevatron inclusive jet data to a prior parton fit using only DIS and DY data. We have seen that when reweighted by the inclusive jet data, this fit becomes statistically equivalent to a refitting using all the data. The statistical equivalence has been quantified using the distance between prior and reweighted sets. This confirms that the refitted and the reweighted PDF sets can be seen as two samples of the same underlying probability distribution. This is simultaneously a validation of the reweighting methodology, and an important a posteriori consistency check of the fitting procedure: an explicit confirmation that reweighting is equivalent to refitting for all data included in the global fit would amount to a proof that the fitted result is indeed that dictated by the laws of statistical inference.

Using the reweighting formalism we have determined the impact of recent high luminosity D0 Run II lepton asymmetry data on the NNPDF2.0 PDFs. The lepton asymmetry data has been historically an important constraint on the large- x d/u ratio, but recent attempts [11,12] to include the new D0 data into global fits have been problematic. We find instead that the data which are inclusive in E_T^l , the muon asymmetry data [8] and electron asymmetry data [9] with $p_T > 25$ GeV, are fully consistent with the NNPDF2.0 predictions and have a moderate impact on PDFs, showing up as a modest though noticeable reduction in the uncertainty of the valence quark distribution. Moreover they are consistent with each other and with all the other datasets included in NNPDF2.0.

The consistency of these data has been recently studied also by the MSTW and CTEQ Collaborations. In particular MSTW [12] finds that it is not possible to fit the inclusive D0 electron dataset without affecting the description of the rest of the experiments in the global analysis unless large nuclear corrections for the DIS deuteron data are applied at the same time. The CT10 analysis [11] also suggests a sizable tension between the D0 lepton asymmetry data and the DIS deuteron data. Our results do not support these conclusions. Since the predictions for the lepton asymmetry depend strongly on the d/u slope, it is possible that the origin of the problems in the

CT10 and MSTW analyses is that they are based on refitting using a fixed parametrization, and are thus subject to the functional biases such a procedure necessarily entails.

We further find that the less inclusive electron asymmetry data [9] binned in E_T^e , the two datasets with $25 \text{ GeV} < E_T^e < 35 \text{ GeV}$ and $35 \text{ GeV} < E_T^e$, while having potentially more impact on the PDFs, are problematic: the former data set is inconsistent with some of the DIS data (specifically BCDMS, both proton and deuteron), while the latter seems to have problems of internal consistency.⁵ Consequently the effect on PDFs of including these datasets in the analysis is to actually increase uncertainties in some regions of x . Furthermore, we find evidence that these two datasets are also mutually inconsistent. We think it likely that the experimental errors on these data have been substantially underestimated. Until these problems are better understood, we believe that is safer to include in the global fit only the inclusive datasets, which even if less constraining are more robust experimentally.

The reweighting methodology described here should allow anybody to perform their own updates of NNPDF fits, to incorporate whatever new datasets they are interested in, by following the same procedure we used here for the specific case of the W lepton asymmetry. We very much hope that they will exploit this possibility.

Acknowledgements

We would like to thank G. Ferrera and M. Grazzini for useful discussions and for instructing us in using the DYNLO code. We are also grateful to G. Cowan and W. Giele for discussions about reweighting and Bayesian statistics, P. Nadolsky for discussions about the W asymmetry data, and R. McNulty and F. De Lorenzi for discussions about the LHCb study. M.U. is supported by the Bundesministerium für Bildung und Forschung (BmBF) of the Federal Republic of Germany (project code 05H09PAE). This work was partly supported by HEPTOOLS under contract MRTN-CT-2006-035505. We would also like to acknowledge the use of the computing resources provided by the Edinburgh Computer and Data Facility (ECDF) (<http://www.ecdf.ed.ac.uk/>). The ECDF is partially supported by the eDIKT initiative (<http://www.edikt.org.uk>).

Appendix A. Distances between reweighted PDFs

Given two sets of $N_{\text{rep}}^{(1)}$ and $N_{\text{rep}}^{(2)}$ replicas, in general reweighted, it is possible to use the distance estimators defined in Appendix A of Ref. [4] to determine whether they correspond to different instances of the same underlying probability distribution, or whether instead they come from different underlying distributions.

The discussion in Ref. [4] applies also to reweighted PDF sets with the corresponding modifications that we list below. For example, expectation values have to be computed with the associated weights. For the first, second and fourth moments of the PDFs one then has to use

$$\langle q^{(k)} \rangle_{(i)} = \frac{1}{N_{\text{rep}}^{(i)}} \sum_{k=1}^{N_{\text{rep}}^{(i)}} w_k^{(i)} q_k^{(i)}, \quad (39)$$

⁵ Similar difficulties in fitting these datasets have been reported by MSTW [12] and CTEQ [11], though it is not easy to make a direct comparison since they attempt to fit all three D0 electron bins simultaneously.

$$\sigma_{(i)}^2[q^{(i)}] = \frac{1}{N_{\text{rep}}^{(i)} - 1} \sum_{k=1}^{N_{\text{rep}}^{(i)}} w_k^{(i)} (q_k^{(i)} - \langle q^{(i)} \rangle)^2, \quad (40)$$

$$m_4[q^{(i)}] = \frac{1}{N_{\text{rep}}^{(i)}} \sum_{k=1}^{N_{\text{rep}}^{(i)}} w_k^{(i)} (q_k^{(i)} - \langle q^{(i)} \rangle)^4. \quad (41)$$

Note that in the above equations the unweighted expressions are trivially reproduced setting $w_k^{(i)} = 1$.

Another difference arises when computing the variance of the mean and the variance of the variance with weighted PDF sets. In this case, these estimators scale not with the total number of replicas but with an effective number of replicas after reweighting

$$N_{\text{rep,eff}}^{(i)} = \frac{(\sum_{k=1}^{N_{\text{rep}}^{(i)}} w_k^{(i)})^2}{\sum_{k=1}^{N_{\text{rep}}^{(i)}} w_k^{(i),2}} = N_{\text{rep}}^2 / \left(\sum_{k=1}^{N_{\text{rep}}^{(i)}} w_k^{(i),2} \right). \quad (42)$$

that reduces to N_{rep} in the unweighted case (note that this is not the same as the N_{eff} given by the Shannon entropy (10)).

The variance of the mean for reweighted sets is then given by

$$\sigma_{(i)}^2[\langle q^{(i)} \rangle] = \frac{1}{N_{\text{rep,eff}}^{(i)}} \sigma_{(i)}^2[q^{(i)}] \quad (43)$$

while the variance of the sample variance is

$$\sigma_{(i)}^2[\bar{\sigma}_{(i)}^2] = \frac{1}{N_{\text{rep,eff}}^{(i)}} \left[m_4[q^{(i)}] - \frac{N_{\text{rep}}^{(i)} - 3}{N_{\text{rep}}^{(i)} - 1} (\bar{\sigma}_{(i)}^2)^2 \right]. \quad (44)$$

Again in the unweighted case everything reduces to the expression in Ref. [4].

References

- [1] NNPDF Collaboration, L. Del Debbio, et al., JHEP 0703 (2007) 039, arXiv:hep-ph/0701127.
- [2] NNPDF Collaboration, R.D. Ball, et al., Nucl. Phys. B 809 (2009) 1, arXiv:0808.1231.
- [3] NNPDF Collaboration, R.D. Ball, et al., Nucl. Phys. B 823 (2009) 195, arXiv:0906.1958.
- [4] NNPDF Collaboration, R.D. Ball, et al., Nucl. Phys. B 838 (2010) 136, arXiv:1002.4407.
- [5] W.T. Giele, S. Keller, Phys. Rev. D 58 (1998) 094023, arXiv:hep-ph/9803393.
- [6] NNPDF Collaboration, L. Del Debbio, et al., JHEP 0503 (2005) 080, arXiv:hep-ph/0501067.
- [7] CDF Collaboration, D.E. Acosta, et al., Phys. Rev. D 71 (2005) 051104, arXiv:hep-ex/0501023.
- [8] D0 Collaboration, V.M. Abazov, et al., Phys. Rev. D 77 (2008) 011106, arXiv:0709.4254.
- [9] D0 Collaboration, V.M. Abazov, et al., Phys. Rev. Lett. 101 (2008) 211801, arXiv:0807.3367.
- [10] M. Vesterinen, D0 Conference note 5976-CONF, arXiv:1006.0451, 2010.
- [11] H.-L. Lai, et al., Phys. Rev. D 82 (2010) 074024, arXiv:1007.2241.
- [12] R.S. Thorne, et al., PoS (2010) 052, arXiv:1006.2753.
- [13] W.T. Giele, D. Kosower, S. Keller, arXiv:hep-ph/0104052.
- [14] F. De Lorenzi, arXiv:1011.4260.
- [15] NNPDF Collaboration, R.D. Ball, et al., JHEP 1005 (2010) 075, arXiv:0912.2276.
- [16] E.T. Jaynes, Probability Theory: The Logic of Science, Cambridge University Press, ISBN 0-521-59271-2, 2003.
- [17] CDF Collaboration, A. Abulencia, et al., Phys. Rev. D 75 (2007) 092006, arXiv:hep-ex/0701051.
- [18] D0 Collaboration, V.M. Abazov, et al., Phys. Rev. Lett. 101 (2008) 062001, arXiv:0802.2400.
- [19] CDF Collaboration, T. Aaltonen, et al., Phys. Rev. Lett. 102 (2009) 181801, arXiv:0901.2169.

- [20] T. Carli, et al., Eur. Phys. J. C 66 (2010) 503, arXiv:0911.2985.
- [21] A.D. Martin, et al., Eur. Phys. J. C 63 (2009) 189, arXiv:0901.0002.
- [22] S. Catani, et al., Phys. Rev. Lett. 103 (2009) 082001, arXiv:0903.2120.
- [23] S. Catani, G. Ferrera, M. Grazzini, JHEP 1005 (2010) 006, arXiv:1002.3115.
- [24] C. Balazs, C.P. Yuan, Phys. Rev. D 56 (1997) 5558, arXiv:hep-ph/9704258.
- [25] M. Dittmar, et al., arXiv:0901.2504.
- [26] J.M. Campbell, R.K. Ellis, Phys. Rev. D 62 (2000) 114012, arXiv:hep-ph/0006304.
- [27] J. Campbell, R.K. Ellis, Phys. Rev. D 65 (2002) 113007, arXiv:hep-ph/0202176.
- [28] MCFM, <http://mcfm.fnal.gov>.