

IPPP/08/95
DCPT/08/190
Cavendish-HEP-08/16
11th June 2009

Parton distributions for the LHC

A.D. MARTIN^a, W.J. STIRLING^b, R.S. THORNE^c AND G. WATT^c

^a Institute for Particle Physics Phenomenology, University of Durham, DH1 3LE, UK

^b Cavendish Laboratory, University of Cambridge, CB3 0HE, UK

^c Department of Physics and Astronomy, University College London, WC1E 6BT, UK

Abstract

We present updated leading-order, next-to-leading order and next-to-next-to-leading order parton distribution functions (“MSTW 2008”) determined from global analysis of hard-scattering data within the standard framework of leading-twist fixed-order collinear factorisation in the $\overline{\text{MS}}$ scheme. These parton distributions supersede the previously available “MRST” sets and should be used for the first LHC data-taking and for the associated theoretical calculations. New data sets fitted include CCFR/NuTeV dimuon cross sections, which constrain the strange quark and antiquark distributions, and Tevatron Run II data on inclusive jet production, the lepton charge asymmetry from W decays and the Z rapidity distribution. Uncertainties are propagated from the experimental errors on the fitted data points using a new dynamic procedure for each eigenvector of the covariance matrix. We discuss the major changes compared to previous MRST fits, briefly compare to parton distributions obtained by other fitting groups, and give predictions for the W and Z total cross sections at the Tevatron and LHC.

Contents

1	Introduction	4
2	Survey of experimental developments	6
3	Overview of theoretical framework	9
4	Treatment of heavy flavours	13
4.1	Fixed flavour number scheme	13
4.2	Zero-mass variable flavour number scheme	14
4.3	General-mass variable flavour number schemes	15
4.4	Intrinsic heavy flavour	25
5	Global parton analyses	27
5.1	Choice of data sets	27
5.2	χ^2 definition	29
5.3	Discussion of fit results	34
6	Error propagation from experimental data points	37
6.1	Review of Hessian method in global PDF analyses	37
6.2	Dynamic determination of tolerance	42
6.3	Uncertainties on input PDFs	47
6.4	Effect of free data set normalisations on PDF uncertainties	53
6.5	Input parameterisation and uncertainties	55
6.6	Fit to a reduced dataset	63
7	Description of inclusive structure function data	69
7.1	HERA structure functions	69
7.2	Fixed-target neutral-current structure functions	72
7.3	Neutrino structure functions	72
8	Neutrino dimuon production: constraints on strangeness	79
8.1	Description of dimuon cross sections	80
8.2	Impact on strange and antistrange distributions	85
8.3	Implications for the NuTeV $\sin^2 \theta_W$ anomaly	87
9	Description of $F_2^{c\bar{c}}$ and $F_2^{b\bar{b}}$ data	89
9.1	Comparison to HERA data	89
9.2	Constraints on intrinsic charm	92
10	Description of low-energy Drell–Yan data	94

11	W and Z boson production at the Tevatron	97
11.1	Description of Tevatron W asymmetry data	97
11.2	Description of Tevatron Z rapidity distribution data	103
11.3	Impact on the down valence quark distribution	106
12	Description of jet data	107
12.1	Description of Run I Tevatron jet data	107
12.2	Description of Run II Tevatron jet data	108
12.3	Impact on the gluon distribution	111
12.4	Jet production at HERA	116
13	Low-x gluon distribution and F_L	122
14	Comparison with other PDF sets	128
15	W and Z total cross section predictions	136
16	Conclusions	144

1 Introduction

In deep-inelastic scattering (DIS), and “hard” proton–proton (or proton–antiproton) high-energy collisions, the scattering proceeds via the partonic constituents of the hadron. To predict the rates of the various processes a set of universal parton distribution functions (PDFs) is required. These distributions are best determined by global fits to all the available DIS and related hard-scattering data. The fits can be performed at leading-order (LO), next-to-leading order (NLO) or at next-to-next-to-leading order (NNLO) in the strong coupling α_S . Over the last couple of years there has been a considerable improvement in the precision, and in the kinematic range, of the experimental measurements for many of these processes, as well as new types of data becoming available. In addition, there have been valuable theoretical developments, which increase the reliability of the global analyses. It is therefore timely, particularly in view of the forthcoming experiments at the Large Hadron Collider (LHC) at CERN, to perform new global analyses which incorporate all of these improvements.

The year 2008 marked the twentieth anniversary of the publication of the first MRS analysis of parton distributions [1], which was, indeed, the first NLO global analysis. Following this initial MRS analysis there have been numerous updates over the years, necessitated by both the regular appearance of new data sets and by theoretical developments [2–21]. In the modern era, the MRST98 (NLO) sets [12] were the first of our sets to take full advantage of a large amount of new HERA structure function data, and the first to incorporate heavy quarks in a consistent and rigorous way in the default PDF set (using a general-mass variable flavour number scheme approach; see Section 4). The uncertainty in the PDFs was explored by producing a modest number of additional sets with different gluon distributions and different values of the strong coupling. The following year, the NLO set was updated (MRST99), a LO set was produced, and new sets corresponding to a different treatment of higher-twist contributions and the DIS scheme were presented [13]. The year 2001 saw a major update [14], with a NNLO (MRST 2001) set, based on an approximation of the NNLO splitting functions [22, 23], produced for the first time [15], alongside new NLO and LO sets. The grid interpolation was also improved to allow for faster and more accurate access to the PDFs in the public interface code. The following year the Hessian approach (see Section 6) was used to produce a “parton distributions with errors” package (MRST 2001 E) comprising a central NLO set and 30 extremum sets [16]. The central NLO and NNLO sets were updated slightly in the same year (MRST 2002) [16], using an improved approximation to the NNLO splitting functions [24]. In 2003, fits were performed in which the x and Q^2 range of DIS structure function data was restricted to ensure stability with respect to cuts on the data, and corresponding NLO and NNLO “conservative” variants of the MRST 2002 sets were derived (MRST 2003 C) [17]. The next major milestone was in 2004, with a substantial update of the NLO and NNLO sets (MRST 2004) [18], the latter using the full NNLO splitting functions [25, 26] for the first time and both incorporating a “physical” parameterisation of the gluon distribution in order to better describe the high- E_T Tevatron jet data. A NLO set incorporating $\mathcal{O}(\alpha)$ QED corrections in the DGLAP evolution equations was also produced for the first time (MRST 2004 QED) [19], together with fixed-flavour-number LO and NLO variants [20]. Finally, in 2006 a NNLO set “with errors” was produced for the first time (MRST 2006 NNLO) [21], using a new general-mass variable flavour number scheme and

MSTW 2008 NLO PDFs (68% C.L.)




Figure 1: MSTW 2008 NLO PDFs at $Q^2 = 10 \text{ GeV}^2$ and $Q^2 = 10^4 \text{ GeV}^2$.

with broader grid coverage in x and Q^2 than in previous sets.

In this paper we present the new MSTW 2008 PDFs at LO, NLO and NNLO. These sets are a major update to the currently available MRST 2001 LO [15], MRST 2004 NLO [18] and MRST 2006 NNLO [21] PDFs. The “end products” of the present paper are grids and interpolation code for the PDFs, which can be found at Ref. [27]. An example is given in Fig. 1, which shows the NLO PDFs at scales of $Q^2 = 10 \text{ GeV}^2$ and $Q^2 = 10^4 \text{ GeV}^2$, including the associated one-sigma (68%) confidence level (C.L.) uncertainty bands.

The contents of this paper are as follows. The new experimental information is summarised in Section 2. An overview of the theoretical framework is presented in Section 3 and the treatment of heavy flavours is explained in Section 4. In Section 5 we present the results of the global fits and in Section 6 we explain the improvements made in the error propagation of the experimental data to the PDF uncertainties, and their consequences. Then we present a more detailed discussion of the description of different data sets included in the global fit: inclusive DIS structure functions (Section 7), dimuon cross sections from neutrino–nucleon scattering (Section 8), heavy flavour DIS structure functions (Section 9), low-energy Drell–Yan production (Section 10), W and Z production at the Tevatron (Section 11), and inclusive jet production at the Tevatron and at HERA (Section 12). In Section 13 we discuss the low- x gluon and the description of the longitudinal structure function, in Section 14 we compare our PDFs with other recent sets, and in Section 15 we present predictions for W and Z total cross sections at the Tevatron and LHC. Finally, we conclude in Section 16. Throughout the text we will highlight the numerous refinements and improvements made to the previous MRST analyses.

2 Survey of experimental developments

Since the most recent MRST analyses [15, 18, 21] a large number of new data sets suitable for inclusion in the global fit have become available, or are included for the first time. Some of these are entirely new types of data, while others supersede existing sets, either improving the precision, extending the kinematic range, or both. Here, we list the new data that we include in the global fit, together with an indication of the parton distributions that they mainly constrain.

- (i) Compared to the analysis in Ref. [18] there is no new large $ep \rightarrow eX$ deep-inelastic F_2 structure function data set. However, note that we now fit to the measured reduced cross section values

$$\tilde{\sigma}(x, Q^2) = F_2(x, Q^2) - \frac{y^2}{1 + (1 - y)^2} F_L(x, Q^2), \quad (1)$$

where $y = Q^2/(xs)$ and \sqrt{s} is the ep centre-of-mass energy. In fact, this was already done in Refs. [21, 28]. Since $y \ll 1$ in most of the kinematic range, $\tilde{\sigma}(x, Q^2)$ is effectively the same as $F_2(x, Q^2)$. However, at HERA, for the lowest x -values at a given Q^2 , the value of y can become as large as 0.7–0.8, and the effect of $F_L(x, Q^2)$ becomes apparent [29, 30]. We now include a small, but important, number of data points from H1 in Ref. [31]. The effect of $F_L(x, Q^2)$ is seen in the data as a flattening of the growth of $\tilde{\sigma}(x, Q^2)$ as x decreases to very small values (for fixed Q^2), leading eventually to a turnover. Hence, for precise analysis of the HERA structure function data it is particularly important to fit any theoretical prediction to the measured $\tilde{\sigma}(x, Q^2)$, rather than to model-dependent extracted values of $F_2(x, Q^2)$. In previous analyses we have used the next best approach to fitting to $\tilde{\sigma}(x, Q^2)$ directly by making the correction using our own prediction for $F_L(x, Q^2)$, maintaining self-consistency. We now also include data on $F_L(x, Q^2)$ at high x from fixed-target experiments [32–34], which provide a weak constraint on the gluon for $x \sim 0.1$.

- (ii) The structure functions F_2 and xF_3 for the charged-current deep-inelastic processes $e^-p \rightarrow \nu_e X$ and $e^+p \rightarrow \bar{\nu}_e X$ have been measured at HERA [35, 36]. At large x these processes are dominated by the u and d valence quark distributions respectively, although the event rate is low in this domain. In principle, these data are valuable as they constrain exactly the same partonic combinations as the data which exist for the neutrino processes $\bar{\nu}_\mu N \rightarrow \mu^+ X$ and $\nu_\mu N \rightarrow \mu^- X$, but without the problems of having to allow for the effects of the heavy nuclear target, N , but the precision of published data is still relatively low. We include the data on $e^+p \rightarrow \bar{\nu}_e X$ since these are more precise, and constrain the less well-known down valence distribution.
- (iii) We include new neutrino structure function data from NuTeV (with an iron target) [37] and CHORUS (with a lead target) [38], though consistency requires the use of only data with $x < 0.5$, as discussed later in Section 7.3. These provide information on the flavour decomposition of quarks for $0.01 \lesssim x \lesssim 0.5$.
- (iv) Data also from the neutrino DIS experiments NuTeV and CCFR on dimuon production give a direct [39] constraint on the strange quark content of the proton for $0.01 \lesssim x \lesssim 0.2$,

replacing the previous direct, but model-dependent, constraint on the strange distribution [40]. Furthermore, the separation into neutrino and antineutrino cross sections allows a separation into quark and antiquark contributions and therefore a determination of the strangeness asymmetry.

- (v) There are now improved measurements of the structure functions, $F_2^{c\bar{c}}$ and $F_2^{b\bar{b}}$ [41–47], for heavy quark production at HERA. We include all of the published data on $F_2^{c\bar{c}}$ in the analysis. Since these HERA data are driven by the small- x gluon distribution, via the $g \rightarrow H\bar{H}$ transition, they probe the gluon for $x \sim 0.001$. However, they also provide useful information on the mass of the charm quark and are a test of our procedure for including heavy flavours. Since the data on $F_2^{b\bar{b}}$ are far less numerous and of lower precision, they do not provide any further constraints on the gluon and so we simply compare to these.
- (vi) There are Tevatron Run II data on the lepton charge asymmetry from W decays [48–50]. These are more precise than the Run I data [51] and provide constraints on the down quark distribution for $x \gtrsim 0.05$, and also to a smaller extent on the up and down sea quarks for $x \approx 0.1$.
- (vii) For the first time there are data from CDF [52] and D0 [53] on the $Z (\rightarrow \ell^+\ell^-)$ rapidity distribution which constrain the quarks for $x \gtrsim 0.05$ with a different weighting than DIS data and at a significantly higher Q^2 scale.
- (viii) There now exist data for inclusive jet production from Run II at the Tevatron from CDF [54, 55] and DØ [56]. These data exist in different rapidity bins, are more precise than the previous Run I data and go out to larger jet p_T values. They provide constraints on the gluon (and quark) distributions in the domain $0.01 \lesssim x \lesssim 0.5$.
- (ix) We include data from HERA on inclusive jet production in DIS [57–59]. (We do not include the photoproduction data because of potential sensitivity to photon PDFs.) These data constrain the gluon for $0.01 \lesssim x \lesssim 0.1$.

The total data that we use in the new global analyses are listed in Table 2 of Section 5, together with the individual χ^2 values for each data set for the LO, NLO and NNLO fits. A rough indication of the particular parton distributions that the various data constrain is also given in Table 1.

Process	Subprocess	Partons	x range
$\ell^\pm \{p, n\} \rightarrow \ell^\pm X$	$\gamma^* q \rightarrow q$	q, \bar{q}, g	$x \gtrsim 0.01$
$\ell^\pm n/p \rightarrow \ell^\pm X$	$\gamma^* d/u \rightarrow d/u$	d/u	$x \gtrsim 0.01$
$pp \rightarrow \mu^+ \mu^- X$	$u\bar{u}, d\bar{d} \rightarrow \gamma^*$	\bar{q}	$0.015 \lesssim x \lesssim 0.35$
$pn/pp \rightarrow \mu^+ \mu^- X$	$(u\bar{d})/(u\bar{u}) \rightarrow \gamma^*$	\bar{d}/\bar{u}	$0.015 \lesssim x \lesssim 0.35$
$\nu(\bar{\nu}) N \rightarrow \mu^-(\mu^+) X$	$W^* q \rightarrow q'$	q, \bar{q}	$0.01 \lesssim x \lesssim 0.5$
$\nu N \rightarrow \mu^- \mu^+ X$	$W^* s \rightarrow c$	s	$0.01 \lesssim x \lesssim 0.2$
$\bar{\nu} N \rightarrow \mu^+ \mu^- X$	$W^* \bar{s} \rightarrow \bar{c}$	\bar{s}	$0.01 \lesssim x \lesssim 0.2$
$e^\pm p \rightarrow e^\pm X$	$\gamma^* q \rightarrow q$	g, q, \bar{q}	$0.0001 \lesssim x \lesssim 0.1$
$e^+ p \rightarrow \bar{\nu} X$	$W^+ \{d, s\} \rightarrow \{u, c\}$	d, s	$x \gtrsim 0.01$
$e^\pm p \rightarrow e^\pm c\bar{c} X$	$\gamma^* c \rightarrow c, \gamma^* g \rightarrow c\bar{c}$	c, g	$0.0001 \lesssim x \lesssim 0.01$
$e^\pm p \rightarrow \text{jet} + X$	$\gamma^* g \rightarrow q\bar{q}$	g	$0.01 \lesssim x \lesssim 0.1$
$p\bar{p} \rightarrow \text{jet} + X$	$gg, qg, q\bar{q} \rightarrow 2j$	g, q	$0.01 \lesssim x \lesssim 0.5$
$p\bar{p} \rightarrow (W^\pm \rightarrow \ell^\pm \nu) X$	$ud \rightarrow W, \bar{u}\bar{d} \rightarrow W$	u, d, \bar{u}, \bar{d}	$x \gtrsim 0.05$
$p\bar{p} \rightarrow (Z \rightarrow \ell^+ \ell^-) X$	$uu, dd \rightarrow Z$	d	$x \gtrsim 0.05$

Table 1: The main processes included in the current global PDF analysis ordered in three groups: fixed-target experiments, HERA and the Tevatron. For each process we give an indication of their dominant partonic subprocesses, the primary partons which are probed and the approximate range of x constrained by the data.

3 Overview of theoretical framework

In this section we first give a brief overview of the standard theoretical formalism used, and then present a summary of the theoretical improvements and changes in methodology in the global analysis. A more detailed discussion of the various items is given later in separate sections.

We work within the standard framework of leading-twist fixed-order collinear factorisation in the $\overline{\text{MS}}$ scheme, where structure functions in DIS, $F_i(x, Q^2)$, can be written as a convolution of coefficient functions, $C_{i,a}$, with PDFs of flavour a in a hadron of type A , $f_{a/A}(x, Q^2)$, i.e.

$$F_i(x, Q^2) = \sum_{a=q,g} C_{i,a} \otimes f_{a/A}(x, Q^2). \quad (2)$$

Similarly, in hadron–hadron collisions, hadronic cross sections can be written as process-dependent partonic cross sections convoluted with the same universal PDFs, i.e.

$$\sigma_{AB} = \sum_{a,b=q,g} \hat{\sigma}_{ab} \otimes f_{a/A}(x_1, Q^2) \otimes f_{b/B}(x_2, Q^2). \quad (3)$$

The scale dependence of the PDFs is given by the DGLAP evolution equation in terms of the calculable splitting functions, $P_{aa'}$, i.e.

$$\frac{\partial f_{a/A}}{\partial \ln Q^2} = \sum_{a'=q,g} P_{aa'} \otimes f_{a'/A}. \quad (4)$$

The DIS coefficient functions, $C_{i,a}$, the partonic cross sections, $\hat{\sigma}_{ab}$, and the splitting functions, $P_{aa'}$, can each be expanded as perturbative series in the running strong coupling, $\alpha_S(Q^2)$. The strong coupling satisfies the renormalisation group equation, which up to NNLO reads

$$\frac{d}{d \ln Q^2} \left(\frac{\alpha_S}{4\pi} \right) = -\beta_0 \left(\frac{\alpha_S}{4\pi} \right)^2 - \beta_1 \left(\frac{\alpha_S}{4\pi} \right)^3 - \beta_2 \left(\frac{\alpha_S}{4\pi} \right)^4 - \dots \quad (5)$$

The input for the evolution equations, (4) and (5), $f_{a/A}(x, Q_0^2)$ and $\alpha_S(Q_0^2)$, at a reference input scale, taken to be $Q_0^2 = 1 \text{ GeV}^2$, must be determined from a global analysis of data. In the present study we use a slightly extended form, compared to previous MRST fits, of the parameterisation of the parton distributions at the input scale $Q_0^2 = 1 \text{ GeV}^2$:

$$xu_v(x, Q_0^2) = A_u x^{\eta_1} (1-x)^{\eta_2} (1 + \epsilon_u \sqrt{x} + \gamma_u x), \quad (6)$$

$$xd_v(x, Q_0^2) = A_d x^{\eta_3} (1-x)^{\eta_4} (1 + \epsilon_d \sqrt{x} + \gamma_d x), \quad (7)$$

$$xS(x, Q_0^2) = A_S x^{\delta_S} (1-x)^{\eta_S} (1 + \epsilon_S \sqrt{x} + \gamma_S x), \quad (8)$$

$$x\Delta(x, Q_0^2) = A_\Delta x^{\eta_\Delta} (1-x)^{\eta_S+2} (1 + \gamma_\Delta x + \delta_\Delta x^2), \quad (9)$$

$$xg(x, Q_0^2) = A_g x^{\delta_g} (1-x)^{\eta_g} (1 + \epsilon_g \sqrt{x} + \gamma_g x) + A_{g'} x^{\delta_{g'}} (1-x)^{\eta_{g'}}, \quad (10)$$

$$x(s + \bar{s})(x, Q_0^2) = A_+ x^{\delta_S} (1-x)^{\eta_+} (1 + \epsilon_S \sqrt{x} + \gamma_S x), \quad (11)$$

$$x(s - \bar{s})(x, Q_0^2) = A_- x^{\delta_-} (1-x)^{\eta_-} (1 - x/x_0), \quad (12)$$

where $\Delta \equiv \bar{d} - \bar{u}$, $q_v \equiv q - \bar{q}$, and where the light quark sea contribution is defined as

$$S \equiv 2(\bar{u} + \bar{d}) + s + \bar{s}. \quad (13)$$

The input PDFs listed in Eqs. (6)–(12) are subject to three constraints from number sum rules:

$$\int_0^1 dx \, u_v(x, Q_0^2) = 2, \quad \int_0^1 dx \, d_v(x, Q_0^2) = 1, \quad \int_0^1 dx \, s_v(x, Q_0^2) = 0, \quad (14)$$

together with the momentum sum rule:

$$\int_0^1 dx \, x [u_v(x, Q_0^2) + d_v(x, Q_0^2) + S(x, Q_0^2) + g(x, Q_0^2)] = 1. \quad (15)$$

We use these four constraints to determine A_g , A_u , A_d and x_0 in terms of the other parameters. There are therefore potentially $34 - 4 = 30$ free PDF parameters in the fit, including α_S . The values of the parameters obtained in the LO, NLO and NNLO fits are given in Table 4 in Section 5 below. (In practice, we fix $\delta_- = 0.2$ in (12) due to extreme correlation with A_- and η_- .) For the LO fit, where there is no tendency for the input gluon distribution to go negative at small x , the second term of the parameterisation (10) is omitted.

The major changes to the analysis framework, compared to previous MRST analyses, are listed below.

- (i) We produce and present PDF sets at LO, NLO and NNLO in α_S . The last of these uses the splitting functions calculated in Refs. [25, 26] and the (massless) coefficient functions for structure functions calculated in Refs. [60–65] together with the massive coefficient functions described in Section 4. We do not include here a set of modified LO PDFs for use in LO Monte Carlo generators of the form described in Refs. [66, 67]. This will be the topic of a further study. We produce eigenvector sets describing the uncertainty due to experimental errors in all cases. This is the first time this has been done at LO. We have significantly modified the method of determining the size of the uncertainties, no longer using a simple $\Delta\chi^2 = 50$ above the global minimum to determine the 90% confidence level uncertainty for the PDFs. We now have a *dynamic* determination of the tolerance for each eigenvector direction, which gives qualitatively similar results to previously. That is, we now have no *a priori* fixed value for the tolerance. This method is described in detail in Section 6. We also include, for the first time, the uncertainties in the PDFs due to the uncertainties in the normalisations of the individual data sets included in the global fit. We present the grids for the eigenvector PDFs as both one-sigma (68%) and 90% confidence level uncertainties, rather than just the latter.
- (ii) We now use the simpler and more conventional definition of α_S defined in Ref. [68], i.e. we solve (5) truncated at the appropriate order, starting from an input value of $\alpha_S(Q_0^2)$. This is particularly important at NNLO where the coupling becomes discontinuous at the heavy flavour matching points. Further discussion of the difference between the MRST and MSTW definitions of α_S will be given elsewhere. The value of $\alpha_S(Q_0^2)$ is now one

of the fit parameters and replaces the Λ_{QCD} parameter used in MRST fits. Since there is more than one definition of Λ_{QCD} in common use, replacing it by $\alpha_S(Q_0^2)$ reduces the potential scope for misuse.

- (iii) The new definition of the coupling is then identical to that used in the PEGASUS [68] (which uses moment space evolution) and HOPPET [69] (which uses x space evolution) programs, therefore we can now test our evolved PDFs by benchmarking against those two programs, which were shown to display excellent agreement with each other [70, 71]. We use our best-fit input parameterisations at $Q_0^2 = 1 \text{ GeV}^2$ in each case, which are more complicated than those in the toy PDFs used for the benchmarks in Refs. [70, 71], and perform trial evolutions at LO, NLO and NNLO up to $Q^2 = 10^4 \text{ GeV}^2$, finding agreement with both PEGASUS and HOPPET to an accuracy of $\mathcal{O}(0.1\%)$ or less in most regions, with discrepancies a little larger only in regions where the PDFs are very small or at extremely low x where extrapolations are used. (We see larger discrepancies with PEGASUS at $x \sim 0.9$, where the default choice of the Mellin-inversion contour is inaccurate when using our input parameterisation.) We intend to return to more comprehensive comparisons in the future, but for the moment are satisfied that any discrepancies between evolution codes are orders of magnitude lower than the uncertainties on the PDFs.
- (iv) For structure function calculations, both neutral and charged current, we use an improved general-mass variable flavour number scheme (GM-VFNS) both at NLO and more particularly at NNLO. This is based on the procedure defined in Ref. [72], though there are some slight modifications described in detail in Section 4. The new procedure is much easier to generalise to higher orders. At NLO there is simply a change in the details, but at NNLO there is a major correction to the transition across the heavy flavour matching points, first implemented in Ref. [21], which demonstrates that NNLO sets before 2006 should be considered out-of-date. No other available NNLO sets which include heavy flavours currently treat the transition across the matching points in such a complete manner. A full GM-VFNS is not currently available for Drell–Yan production of virtual photons, or W and Z bosons, or for inclusive jet production. For the low-mass Drell–Yan data in our fit, heavy flavours contribute $\ll 1\%$ of the total, so the inaccuracy invoked by the approximation of using the zero-mass scheme is negligible. All other processes are at scales such that the charm mass m_c is effectively very small, and the zero-mass scheme is a very good approximation. The approximation should still be reliable for processes induced by bottom quarks, and in this case the relative contribution is small.
- (v) Vector boson production data can now be described in a fully exclusive way at NNLO. We use the FEWZ code [73], and compare with the results obtained using the RESBOS code [74] which includes NLO+NNLL p_T^W -resummation effects. This allows, for example, a detailed description of the W asymmetry data accounting for the W width (Γ_W) and lepton decay effects.
- (vi) We now implement FASTNLO [75], based on NLOJET++ [76, 77], which allows the inclusion of the NLO hard cross section corrections to both the Tevatron and HERA jet data in the

fitting routine. This improvement replaces the K -factors and pseudo-gluon data that were previously necessary to speed up the fitting procedure.

- (vii) We provide [27] a new form of the grids that list the parton distributions over an extended Q^2 range and which are more dense in x . The extended public grids allows an improved description around the heavy quark thresholds, where at NNLO discontinuities in parton distributions appear. (We also make improvements to the internal grids used in our evolution code, so that the heavy quark thresholds lie exactly on the grid points, which was not the case in the MRST analyses.) Moreover, the grids now contain more parton flavours: $s \neq \bar{s}$ at all orders due to the (mild) evidence from the dimuon data, and additionally $c \neq \bar{c}$ and $b \neq \bar{b}$ at NNLO, due to automatic generation of a small $\mathcal{O}(\alpha_S^2)$ quark–antiquark difference during the NNLO evolution, so at this order the H distribution is (very) slightly different from \overline{H} [78]. (This small asymmetry at NNLO was omitted in previous MRST analyses.)
- (viii) Two features of the parameterisation are worth emphasising. First, as since Ref. [14], we allow the gluon to have a very general form at low x , where it is by far the dominant parton distribution. We will discuss the consequences of this in Section 6.5. Second, we parameterise $s \pm \bar{s}$ as functions of x , rather than assuming, as previously,

$$s = \bar{s} = \frac{\kappa}{2}(\bar{u} + \bar{d}) \quad (16)$$

at the input scale, where κ was a constant fixed by the neutrino dimuon production ($\kappa \sim 0.4\text{--}0.5$); see Section 8. Implicit in our parameterisation is the assumption that the strange sea will have approximately the same shape as the up and down sea at small x . This extra parametric freedom of the gluon and strange quark distributions means that all partons are less constrained, particularly at low x . This, in turn, leads to a more realistic estimate of the uncertainties on all parton distributions.

- (ix) We fix the heavy quark masses at $m_c = 1.40$ GeV and $m_b = 4.75$ GeV, changed from the MRST default values of $m_c = 1.43$ GeV and $m_b = 4.30$ GeV. Our value of $m_b = 4.75$ GeV is close to the calculated $\overline{\text{MS}}$ mass transformed to the pole mass value, using the three-loop relation between the pole and $\overline{\text{MS}}$ masses, of $m_b = 4.800$ GeV [79, 80], while our value of $m_c = 1.40$ GeV is smaller than the calculated pole mass value of $m_c = 1.666$ GeV [79, 80]. If allowed to go free in the global fits, the best-fit values are $m_c = 1.39$ GeV at NLO and $m_c = 1.27$ GeV at NNLO. In a subsequent paper we will present a more detailed discussion of the sensitivity of different data sets to the charm quark mass and present the best-fit values including a determination of the uncertainty. We allow a maximum of five flavours in the evolution and do not include top quarks.

4 Treatment of heavy flavours

The correct treatment of heavy flavours in an analysis of parton distributions is essential for precision measurements at hadron colliders. For example, the cross section for W production at the LHC depends crucially on precise knowledge of the charm quark distribution. Moreover, a correct method of fitting the heavy flavour contribution to structure functions is important because of the knock-on effect it can have on all parton distributions. However, it has become clear in recent years that it is a delicate issue to obtain a proper treatment of heavy flavours. There are various choices that can be made, and also many ways in which subtle mistakes can occur. Both the choices and the mistakes can lead to changes in parton distributions which may be similar to, or even greater than, the quoted uncertainties — though the mistakes usually lead to the more dramatic changes. Hence, we will here provide a full description of our procedure, along with a comparison to alternatives and some illustrations of pitfalls which must be avoided.

First we describe the two distinct regimes for describing heavy quarks where the pictures are relatively simple. These are the so-called fixed flavour number scheme (FFNS) and zero-mass variable flavour number scheme (ZM-VFNS); see Fig. 2.

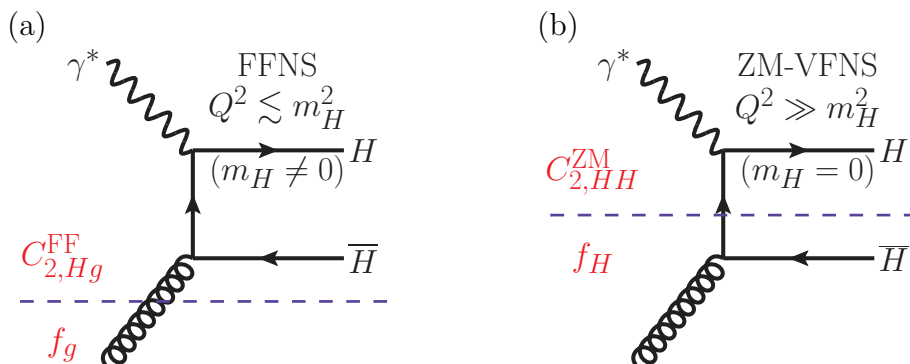


Figure 2: Diagrams contributing at leading-order to the heavy flavour structure function, F_2^H , in (a) the fixed flavour number scheme (FFNS) valid for $Q^2 \lesssim m_H^2$, and (b) the zero-mass variable flavour number scheme (ZM-VFNS) valid for $Q^2 \gg m_H^2$.

4.1 Fixed flavour number scheme

First, there is the region where the hard scale of the process is similar to, or smaller than, the quark mass¹, i.e. $Q^2 \lesssim m_H^2$. In this case it is most natural to describe the massive quarks as final-state particles, and not as partons within the proton. This requirement defines the FFNS, where only light quarks are partons, and the number of flavours is fixed. We label the number of quark flavours appearing as parton distributions by n_f . Note, however, that there are multiple instances of FFNSs with different numbers of *active* flavours n_f . The number n_f is normally equal to 3, where up, down and strange are the light quarks, but we can treat charm as a light

¹Throughout this section we use H to denote a heavy quark: $H = c, b$ or t .

quark while bottom remains heavy at high scales, i.e. $n_f = 4$, and it may also be equal to 5 if only top quarks are treated as heavy. In each example of a FFNS the structure functions are given by²

$$F_i(x, Q^2) = \sum_k C_{i,k}^{\text{FF}, n_f}(Q^2/m_H^2) \otimes f_k^{n_f}(Q^2). \quad (17)$$

This approach contains all the m_H -dependent contributions, and as it is conceptually simple it is frequently used in analyses of structure functions. Even in this case one must be careful to be self-consistent in defining all quantities, i.e. parton distributions, coefficient functions and coupling constant in the same renormalisation schemes (which is often not done). The mistake made by not doing so can lead to errors in the gluon distribution of order the size of the uncertainty [20].

Despite its conceptual simplicity, the FFNS has some potential problems. It does not sum $\alpha_S^m \ln^l(Q^2/m_H^2)$ ($l \leq m$) terms in the perturbative expansion. Thus the accuracy of the fixed-order expansion becomes increasingly uncertain as Q^2 increases above m_H^2 . For example, $\alpha_S(M_W^2) \ln(M_W^2/m_c^2)$ is approximately equal to unity, and so there is no guarantee an expansion in this variable will converge. As well as this problem of principle, there are additional practical issues. Since calculations including full mass dependence are complicated, there are only a few cross sections known even to NLO in α_S within this framework, so the resulting parton distributions are not universally useful. In addition, even for neutral-current structure functions, the FFNS coefficient functions are known only up to NLO [81, 82], and are not calculated at NNLO — that is, the α_S^3 coefficient³, $C_{2,Hg}^{\text{FF}, n_f, (3)}$, for F_2 is unknown, so one cannot determine parton distributions at NNLO in this scheme.

4.2 Zero-mass variable flavour number scheme

All of the problems of the FFNS are solved in the so-called zero-mass variable flavour number scheme (ZM-VFNS). Here, the heavy quarks evolve according to the splitting functions for massless quarks and the resummation of the large logarithms in Q^2/m_H^2 is achieved by the introduction of heavy-flavour parton distributions and the solution of the evolution equations. It assumes that at high scales, $Q^2 \gg m_H^2$, the massive quarks behave like massless partons, and the coefficient functions are simply those in the massless limit, e.g. for structure functions

$$F_i(x, Q^2) = \sum_k C_{i,j}^{\text{ZM}, n_f} \otimes f_j^{n_f}(Q^2), \quad (18)$$

where $n_f - 3$ is the number of active heavy quarks, with masses above some transition point for turning on the heavy flavour distribution, typically at a scale similar to m_H^2 . This is technically simpler than the FFNS, and many more cross sections are known in this scheme. The nomenclature of “zero-mass” is a little misleading because some mass dependence is included in the

²To simplify the discussion, we will use the convention that the factorisation scale and renormalisation scale are both equal to Q^2 . It is also the choice we make in the analysis. Alternative choices are possible, and cause no problems in principle, but can add considerable technical complications.

³We add a subscript H to distinguish the $g \rightarrow H$ coefficient function $C_{2,Hg}$ from the usual coefficient function $C_{2,g}$ describing the $g \rightarrow q$ transition of light quarks.

boundary conditions for evolution. The parton distributions in different quark number regimes are related to each other perturbatively, i.e.

$$f_j^{n+1}(Q^2) = \sum_k A_{jk}(Q^2/m_H^2) \otimes f_k^n(Q^2), \quad (19)$$

where the perturbative matrix elements $A_{jk}(Q^2/m_H^2)$ [83] containing $\ln(Q^2/m_H^2)$ terms are known to NNLO, i.e. $\mathcal{O}(\alpha_S^2)$.⁴ They relate $f_i^n(Q^2)$ and $f_i^{n+1}(Q^2)$, guaranteeing the correct evolution for both regimes. At NLO in the $\overline{\text{MS}}$ scheme they constrain the heavy quarks to begin their evolution from a zero value at $Q^2 = m_H^2$, and the other partons to be continuous at this choice of transition point, hence making it the natural choice.

4.3 General-mass variable flavour number schemes

The ZM-VFNS has many advantages. However, it has the failing that it simply ignores $\mathcal{O}(m_H^2/Q^2)$ corrections to the coefficient functions, and hence it is inaccurate in the region where Q^2 is not so much greater than m_H^2 . The nomenclature *scheme* may be thought of as misleading in a similar way that *zero-mass* might. Scheme usually implies an alternative choice in ordering the expansion, or a particular separation of contributions between coefficient functions and parton distributions, i.e. the inherent ambiguity in a perturbative QCD calculation allows a choice, the effects of which become increasingly smaller as higher orders are included. The ZM-VFNS misses out $\mathcal{O}(m_H^2/Q^2)$ contributions completely, and there is a permanent error of this order. This clearly already happens at LO and NLO. The error induced by fitting to HERA structure function data using a ZM-VFNS was shown to be up to $\sim 6\%$ in the small- x light-quark distributions by CTEQ in Ref. [84], resulting in a systematic error of $\sim 8\%$ in predictions for vector boson production at the LHC. This NLO result seems to provide ample evidence for the use of a general-mass variable flavour number scheme (GM-VFNS) to provide default parton distributions, and indeed this is the approach we have adopted in previous global analyses [12–14, 18, 21]. A definition of a GM-VFNS was first proposed in Ref. [85] — the ACOT scheme, and the MRST group have been using the alternative TR scheme [86] as the default since MRST98 [12]. The CTEQ group have since also adopted this convention [84, 87].

However, the inherent problems with the ZM-VFNS are thrown into particularly sharp relief at NNLO. At $\mathcal{O}(\alpha_S^2)$ the $A_{jk}(Q^2/m_H^2)$ are no longer zero at $Q^2 = m_H^2$, and lead to discontinuities in the parton distributions. This is similar to the discontinuity at $\mu^2 = m_H^2$ in $\alpha_S(\mu^2)$ at NNLO. For the coupling constant this discontinuity is rather small. However, the corresponding discontinuities in parton distributions can be significant, about 10% for the gluon, but it turns out that $c(x, m_c^2)$ is quite considerably negative at small x , see e.g. Fig. 3 of Ref. [21]. As it happens, the effect of the NNLO massless coefficient function is to make this worse rather than better at $x \sim 0.0001 - 0.01$, and the charm contribution to $F_2(x, Q^2)$ is negative at the transition point to such a degree that there is an $\mathcal{O}(10\%)$ discontinuity in the total structure function at $Q^2 = m_c^2$, as illustrated in Fig. 1 of Ref. [72].

⁴When testing our NNLO evolution code against PEGASUS [68] we traced the source of a discrepancy in our implementation of A_{Hg} to a typo in Eq.(B.3) of the preprint version of Ref. [83] (fixed in the journal version). Specifically, the term $C_A T_f (16z + 16z^2) [\text{Li}_2(-z) + \ln z \ln(1+z)]$ has 16 rather than $16z$ in the preprint version.

Hence, for a precise analysis of structure functions, and other data, one must use a GM-VFNS which smoothly connects the two well-defined limits of $Q^2 \leq m_H^2$ and $Q^2 \gg m_H^2$. This is already the case at lower orders, but becomes imperative at NNLO. However, despite the above reasoning, a GM-VFNS is not always used as default even today. This is undoubtedly often because of the extra complication compared to the simple ZM-VFNS. However, part of the reason may be because the definition is not unique. The reasons for this and the consequences will now be discussed, along with a detailed outline of the prescription used in our analysis. A less detailed, but more introductory and comparative discussion of the schemes used by MRST/MSTW and CTEQ can be found in Ref. [88].

A GM-VFNS can be defined by demanding equivalence of the $n_f = n$ (FFNS) and $n_f = n+1$ -flavour (GM-VFNS) descriptions above the transition point for the new parton distributions (they are by definition identical below this point), at all orders, i.e.

$$\begin{aligned} F_i(x, Q^2) &= \sum_k C_{i,k}^{\text{FF},n}(Q^2/m_H^2) \otimes f_k^n(Q^2) = \sum_j C_{i,j}^{\text{VF},n+1}(Q^2/m_H^2) \otimes f_j^{n+1}(Q^2) \\ &\equiv \sum_{j,k} C_{i,j}^{\text{VF},n+1}(Q^2/m_H^2) \otimes A_{jk}(Q^2/m_H^2) \otimes f_k^n(Q^2). \end{aligned} \quad (20)$$

The description where the number of active partons is taken to be $n_f = n$ must be identical to that when it increases, i.e. $n_f = n+1$. Hence, the GM-VFNS coefficient functions satisfy⁵

$$C_{i,k}^{\text{FF},n}(Q^2/m_H^2) = \sum_j C_{i,j}^{\text{VF},n+1}(Q^2/m_H^2) \otimes A_{jk}(Q^2/m_H^2), \quad (21)$$

which, for example, at $\mathcal{O}(\alpha_S)$ gives for $F_2(x, Q^2)$:

$$C_{2,g}^{\text{FF},n,(1)}(Q^2/m_H^2) = C_{2,HH}^{\text{VF},n+1,(0)}(Q^2/m_H^2) \otimes P_{qg}^{(0)} \ln(Q^2/m_H^2) + C_{2,g}^{\text{VF},n+1,(1)}(Q^2/m_H^2). \quad (22)$$

The GM-VFNS coefficient functions, $C_{i,j}^{\text{VF},n_f}$, are constrained to tend to the massless limits for $Q^2 \gg m_H^2$ and the $A_{jk}(Q^2/m_H^2)$ are such that this happens self-consistently. However, the $C_{i,j}^{\text{VF},n_f,(m)}(Q^2/m_H^2)$ are only uniquely defined in this massless limit $Q^2/m_H^2 \rightarrow \infty$. For finite Q^2/m_H^2 one can swap $\mathcal{O}(m_H^2/Q^2)$ terms between $C_{2,HH}^{\text{VF},n+1,(0)}(Q^2/m_H^2)$ and $C_{2,g}^{\text{VF},n+1,(1)}(Q^2/m_H^2)$ while maintaining the exact definition in (22). It is clear that this general feature applies to all relationships in (21). Although the equivalence (21) was first pointed out in general in Ref. [83], and (22) is effectively used in defining the original ACOT scheme, the freedom to swap $\mathcal{O}(m_H^2/Q^2)$ terms without violating the definition of a GM-VFNS was first noticed in Ref. [86] and put to use to define the TR scheme, as described below. This freedom to redistribute $\mathcal{O}(m_H^2/Q^2)$ terms can be classified as a change in *scheme* since it leads to an ambiguity in the result at a fixed order, but the ambiguity becomes higher order if the order of the calculation increases, much like the renormalisation and factorisation scheme (and scale) ambiguities. Moreover, it is a change of scheme which does not change the definition of the parton distributions,

⁵It is implicit that the coupling constant is a function of n flavours on the left-hand side and of $n+1$ flavours on the right-hand side.

only the coefficient functions. This is perhaps a surprising result, which occurs because there is a redundancy in (21), there being one more coefficient function above the transition point than below, i.e. that for the heavy quark.

The original ACOT prescription [85] calculated the coefficient functions for single heavy quark scattering from a virtual photon exactly. This might seem the most natural definition. However, it assumes that immediately above the transition point a *single* heavy quark or antiquark can exist in isolation. Hence, each coefficient function violates the real physical threshold $W^2 > 4m_H^2$, since only *one* heavy quark is produced in the final state, rather than a quark–antiquark *pair*. Moreover, this definition requires the calculation of mass-dependent coefficient functions, which becomes progressively more difficult at higher orders. As mentioned above, in the TR scheme [86] the ambiguity in the definition of $C_{2,HH}^{\text{VF},n_f,(0)}(Q^2/m_H^2)$ was recognised and exploited for the first time. To be precise, the constraint that $(\partial F_2^H/\partial \ln Q^2)$ was continuous at the transition point (in the gluon sector) was applied to define the heavy-quark coefficient functions. This imposed the correct threshold dependence on all coefficient functions and improved the smoothness at $Q^2 = m_H^2$, and did not involve the explicit calculation of new mass-dependent diagrams. However, it did involve complications, since it required the convolution of the formal inverse of splitting functions with coefficient functions, which itself becomes technically difficult at higher orders.

Since these early definitions there have been various modifications, including a precise definition of an ACOT-like scheme up to NNLO by Chuvakin, Smith and van Neerven [89].⁶ A major simplification was achieved when the flexibility in the choice of heavy-quark coefficient functions was used to define the ACOT(χ) prescription [90, 91], which in the language used in this paper (and in Ref. [86]) would be defined by

$$C_{2,HH}^{\text{VF},n_f,(0)}(z, Q^2/m_H^2) = e_H^2 z(1 + 4m_H^2/Q^2) \delta\left(z - \frac{Q^2}{Q^2 + 4m_H^2}\right). \quad (23)$$

This gives the LO definition

$$F_2^{H,(0)}(x, Q^2) = e_H^2 (x/x_{\text{max}})(H + \overline{H})(x/x_{\text{max}}, Q^2), \quad (24)$$

where $x_{\text{max}} = Q^2/(Q^2 + 4m_H^2)$. It automatically reduces to the massless limit $C_{2,HH}^{\text{ZM},(0)}(z) = e_H^2 z \delta(1 - z)$ for $Q^2/m_H^2 \rightarrow \infty$, and also imposes the true physical threshold

$$W^2 = Q^2(1 - x)/x \geq 4m_H^2. \quad (25)$$

⁶Indeed in Ref. [89] it is pointed out that at NNLO a further complication appears for the first time. There are $\ln^3(Q^2/m_H^2)$ divergences at $\mathcal{O}(\alpha_S^2)$ coming from gluon splitting into heavy quark–antiquark pairs. These divergences arise from heavy quark emission diagrams which cancel with opposite ones originating from virtual heavy quark loops in the “light quark” coefficient functions. This cancellation is achieved in a physically meaningful manner by imposing a cut on the softness of the heavy quark final state in the former process, i.e. the unobservable soft process cancels with the virtual corrections in the “light quark” cross section, and a well-behaved remainder enters the “heavy quark” cross section. In practice, after cancellation both contributions are very small (both are quark- rather than gluon-initiated) at NNLO and we currently include the total in the “light quark” sector. In the $\mathcal{O}(\alpha_S^2)$ FFNS they are usually combined instead in the “heavy quark” contribution.

This choice of the LO heavy-flavour coefficient function has been adopted in our current prescription, which we denote the TR' scheme, described in detail in Ref. [72]. For the GM-VFNS to remain simple (and physical) at all orders, m , it is necessary to choose

$$C_{2,HH}^{\text{VF},n_f,(m)}(z, Q^2/m_H^2) = C_{2,HH}^{\text{ZM},n_f,(m)}(z/x_{\text{max}}), \quad (26)$$

which is the implicit ACOT(χ) definition, and is our choice. It removes one of the sources of ambiguity in defining a GM-VFNS. However, there are others.

One major issue in a complete definition of the GM-VFNS, is that of the ordering of the perturbative expansion. This ambiguity comes about because the ordering in α_S for $F_2^H(x, Q^2)$ is different for the number of active flavours $n_f = n$ and $n_f = n + 1$ regions:

n – flavour	$n + 1$ – flavour
LO NLO NNLO	$C_{2,HH}^{\text{VF},n+1,(0)} \otimes (H + \overline{H})$ $\left(C_{2,Hg}^{\text{FF},n,(2)} \otimes g^n + C_{2,Hq}^{\text{FF},n,(2)} \otimes \Sigma^n \right)$ $\left(\frac{\alpha_S}{4\pi} \right)^2 \sum_j C_{2,Hj}^{\text{FF},n,(3)} \otimes f_j^n$
	$C_{2,HH}^{\text{VF},n+1,(1)} \otimes (H + \overline{H}) + C_{2,Hg}^{\text{VF},n+1,(1)} \otimes g^{n+1}$ $\left(\frac{\alpha_S}{4\pi} \right)^2 \sum_j C_{2,Hj}^{\text{VF},n+1,(2)} \otimes f_j^{n+1}$

with obvious generalisation to even higher orders. This means that switching directly from an n flavours fixed order to the $n + 1$ fixed order leads to a discontinuity in $F_2^H(x, Q^2)$. As with the discontinuities in the ZM-VFNS already discussed this is not just a problem in principle — the discontinuity is comparable to the errors on data, particularly at small x .

Hence, any definition of a GM-VFNS must make some decision how to deal with this, and the ACOT-type schemes have always made a different choice to that for the TR-type schemes used in our analyses. The ACOT-type schemes simply define the same order of α_S both below and above the transition point. For example at NLO the definition is

$$F_2^H(x, Q^2) = \frac{\alpha_S}{4\pi} C_{2,Hg}^{\text{FF},n,(1)} \otimes g^n \rightarrow \frac{\alpha_S}{4\pi} \left(C_{2,HH}^{\text{VF},n+1,(1)} \otimes (H + \overline{H}) + C_{2,Hg}^{\text{VF},n+1,(1)} \otimes g^{n+1} \right). \quad (27)$$

This clearly maintains continuity in the structure function across the transition point. However, it only contains information on LO heavy flavour evolution below $Q^2 = m_H^2$, since $C_{2,Hg}^{\text{FF},n,(1)}$ only contains information on the LO splitting function, but the heavy quarks evolve using NLO splitting functions above $Q^2 = m_H^2$ — a big change at small x . The TR scheme, defined in Ref. [86], and all subsequent variations used in our analyses, try to maintain the correct ordering in each region as closely as possible. For example at LO we have the definition

$$\begin{aligned} F_2^H(x, Q^2) &= \frac{\alpha_S(Q^2)}{4\pi} C_{2,Hg}^{\text{FF},n,(1)}(Q^2/m_H^2) \otimes g^n(Q^2) \\ &\rightarrow \frac{\alpha_S(m_H^2)}{4\pi} C_{2,Hg}^{\text{FF},n,(1)}(1) \otimes g^n(m_H^2) + C_{2,HH}^{\text{VF},n+1,(0)}(Q^2/m_H^2) \otimes (H + \overline{H})(Q^2), \end{aligned} \quad (28)$$

i.e. we freeze the $\mathcal{O}(\alpha_S)$ term when going upwards through $Q^2 = m_H^2$. This generalises to higher orders by freezing the term with the highest power of α_S in the definition for $Q^2 < m_H^2$ when

moving upwards above m_H^2 . Hence, the definition of the ordering is consistent within each region, except for the addition of a constant term (which does not affect evolution) above $Q^2 = m_H^2$, which becomes progressively less important at higher Q^2 , and whose power of α_S increases as the order of the perturbative expansion increases.

This definition of the ordering means that in order to define our VFNS at NNLO [72] we need to use the $\mathcal{O}(\alpha_S^3)$ heavy-flavour coefficient functions for $Q^2 \leq m_H^2$ (and this contribution will be frozen for $Q^2 > m_H^2$). As mentioned above, these coefficient functions are not yet calculated. However, as explained in Ref. [72], we can model this contribution using the known leading threshold logarithms [92] and leading $\ln(1/x)$ terms derived from the k_T -dependent impact factors [93]. This results in a significant contribution at small Q^2 and x with some model dependence. However, variation in the free parameters does not lead to a large change, as discussed in Section 9.⁷ Up to this small model dependence we have a full NNLO GM-VFNS with automatic continuity of structure functions across heavy flavour transition points.⁸ This is certainly the most complete treatment of heavy-flavour effects currently used in any NNLO analysis.

Scheme dependence

Although all definitions of the GM-VFNS become very similar at very high Q^2 , the difference in choice can be phenomenologically important. For example, our definition effectively includes exactly one higher order than ACOT-type schemes for $Q^2 < m_H^2$, and the value of this contribution at $Q^2 = m_H^2$ is carried to higher Q^2 . Since at small x and near threshold the higher orders in α_S are accompanied by large corrections, this leads to large differences below the transition point, which are still important a little way above the transition point. This is shown for $F_2^{c\bar{c}}(x, Q^2)$ in Fig. 2 of Ref. [72] where the two choices are shown for the same parton distributions. More clearly, this difference in the definition of the ordering is the main difference in the NLO predictions from MRST and CTEQ in the comparison to H1 data on $F_2^{bb}(x, Q^2)$ [43], shown in Fig. 4 of the same paper [72].

The inclusion of the complete GM-VFNS in a global fit at NNLO first appeared in Ref. [21], and led to some important changes compared to our previous NNLO analyses, which had a much more approximate inclusion of heavy flavours (which was explained clearly in the Appendix of [15]). A consequence of including the positive $\mathcal{O}(\alpha_S^3)$ coefficient functions at low Q^2 is that the NNLO $F_2^{c\bar{c}}(x, Q^2)$ automatically starts from a higher value at low Q^2 . However, at high Q^2 , the structure function is dominated by $(c + \bar{c})(x, Q^2)$. This has started evolving from a

⁷It should be stressed that this model is only valid for the region $Q^2 \leq m_H^2$, and would not be useful for a full NNLO FFNS since it contains no information on the large Q^2/m_H^2 limits of the coefficient functions. A more general approximation to the $\mathcal{O}(\alpha_S^3)$ coefficient functions could be attempted, but full details would require first the calculation of the $\mathcal{O}(\alpha_S^3)$ matrix element A_{Hg} . This more tractable project is being investigated at present [94]. An approximation to the α_S^3 coefficient functions using logarithmically-enhanced terms near threshold and exact scale-dependent terms has recently been proposed in Ref. [95].

⁸There are actually $\mathcal{O}(\alpha_S^3)$ discontinuities due to terms such as $C_{2,HH}^{\text{VF},n+1,(1)} \otimes (H + \overline{H})$ and $C_{2,Hg}^{\text{VF},n+1,(1)} \otimes g^{n+1}$, i.e. $\mathcal{O}(\alpha_S)$ coefficient functions convoluted with $\mathcal{O}(\alpha_S^2)$ discontinuities in partons. These would be cancelled at NNNLO by discontinuities in $\mathcal{O}(\alpha_S^3)$ coefficient functions. In practice the imposition of the correct threshold behaviour in all coefficient functions minimises these effects and they are very small.

significantly negative value at $Q^2 = m_c^2$. The parton distributions in an NNLO fit readjust so that the light flavours evolve similarly to those at NLO, in order to fit the data. Since the heavy flavour quarks evolve at the same rate as light quarks, but at NNLO start from a negative starting value, they remain lower than at NLO for higher Q^2 . Hence, there is a general trend: $F_2^{cc}(x, Q^2)$ is flatter in Q^2 at NNLO than at NLO, as shown in Fig. 4 of Ref. [21]. It is also flatter than in our earlier (approximate) NNLO analyses. We found that this had an important effect on the gluon distribution when we updated our NNLO analysis. As seen in Fig. 5 of Ref. [21], it led to a larger gluon for $x \sim 0.0001$ – 0.01 , as well as a larger value of $\alpha_S(M_Z^2)$, both compensating for the naturally flatter evolution, and consequently leading to more evolution of the light quark sea. Both the gluon and the light quark sea were up to 6–7% greater than in the 2004 set [18] for $Q^2 = 10^4 \text{ GeV}^2$, the increase maximising at $x = 0.0001$ – 0.001 . As a result there was a 6% increase in the predictions for σ_W and σ_Z at the LHC. This surprisingly large change is a *correction* rather than a reflection of the uncertainty due to the freedom in choosing heavy flavour schemes. The treatment of heavy flavour at NNLO is the same in the sets presented in this paper as for the MRST 2006 set, and as we will see, there are no further changes in predictions of the same size as this increase in going from the 2004 to the 2006 analyses. This demonstrates that the MRST 2004 NNLO distributions should now be considered to be obsolete.

Our 2006 NNLO parton update [21] was made because this was the first time the heavy flavour prescription had been treated precisely at NNLO and also because there was previously no MRST NNLO set with uncertainties. The data used in the analysis were very similar to the 2004 set, and since a consistent GM-VFNS was already used at NLO, and a set with uncertainties already existed, no new corresponding release of a NLO set was made along with the 2006 NNLO set. With the benefit of hindsight, it is interesting to check the effect on the distributions due to the change in the prescription for the GM-VFNS at NLO without complicating the issue by also changing many other things in the analysis. To this end we have obtained an unpublished MRST 2006 NLO set, which is fit to exactly the same data as the MRST 2006 NNLO set.⁹

The comparison of the up quark and gluon distributions for the MRST 2006 NLO set and the MRST 2004 NLO set, i.e. the comparable plot to Fig. 5 of Ref. [21] for NNLO, is shown in Fig. 3. As can be seen it leads to the same trend for the parton distributions as at NNLO, i.e. an increase in the small- x gluon and light quarks, but the effect is much smaller — a maximum of a 2% change. Also, the value of the coupling constant increases by 0.001 from the 2004 value of $\alpha_S(M_Z^2) = 0.120$. Again, this is similar to, but smaller than, the change at NNLO. Hence, we can conclude that the change in our choice of the heavy-flavour coefficient function alone leads to changes in the distributions of up to 2%, and since the change is simply a freedom we have in making a definition, this is a theoretical uncertainty on the parton distributions, much like the frequently invoked scale uncertainty. Like the latter, it should decrease as we go to higher orders. The ambiguity simultaneously moves to higher order, but it is difficult to check this explicitly since our main reason for making our change in the choice of heavy-quark coefficient functions was the difficulty of applying the original procedure in Ref. [86] at NNLO. Certainly an absolute maximum of 2% of the 6–7% change, in the predictions for σ_W and σ_Z at the LHC in going from

⁹We do not intend to officially release the MRST 2006 NLO set, since it is superseded by the present MSTW 2008 NLO analysis, but it is nevertheless available on request.

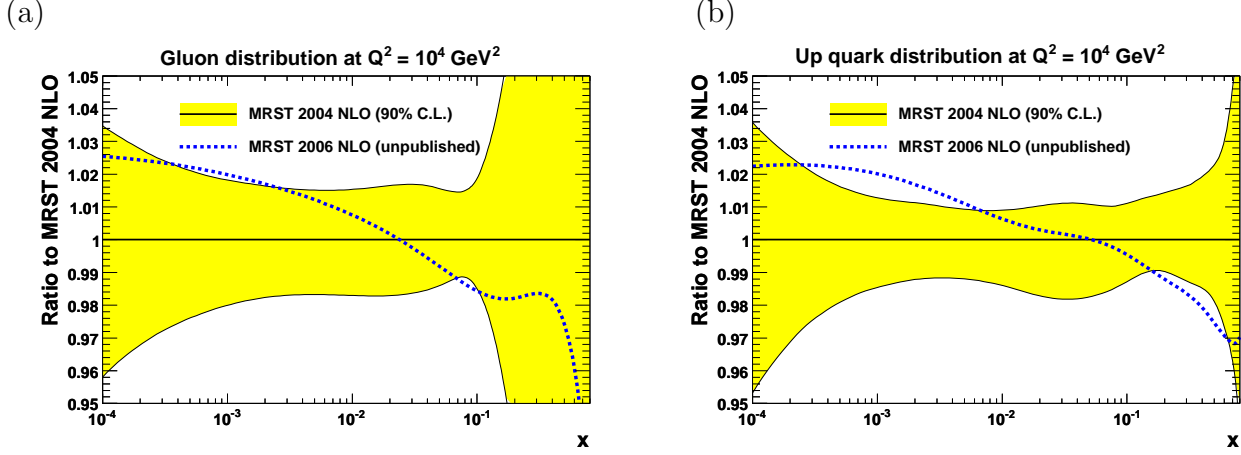


Figure 3: A comparison of the unpublished MRST 2006 NLO parton distributions to the MRST 2004 NLO distributions at $Q^2 = 10^4 \text{ GeV}^2$. In order to illustrate the significance of the size of the differences, the fractional uncertainty on the MRST 2001 distributions is used for the 2004 distributions. The corresponding comparison at NNLO can be seen in Fig. 5 of Ref. [21].

the 2004 to the 2006 NNLO parton sets, is due to true ambiguities, and the remaining 5% is due to the correction of the flaws in the previous approach.

To close this section, we note that the MRST 2006 NLO distributions lead to a nearly 3% increase in the predictions for σ_W and σ_Z at the LHC compared to MRST 2004 NLO, though there is very little change at the Tevatron, where the typical values of x probed are nearly an order of magnitude higher. As with the distributions themselves, this variation is a similar size to the quoted uncertainty in the cross sections, and again this is a genuine theoretical uncertainty. Our most up-to-date predictions for LHC and Tevatron W and Z total cross sections, and comparison to the various previous sets, will be made in Section 15.

Longitudinal structure function

One also has to be careful in defining the GM-VFNS for the longitudinal structure function. In one sense this is not so important since the contribution from the longitudinal structure function to the total reduced cross section measured in DIS experiments is small, and the errors on direct measurements of F_L are comparatively large. However, the importance is increased by the larger ambiguity inherent in the definition of F_L as compared to F_2 .

This large ambiguity occurs because if one calculates the coefficient function for a single massive quark scattering off a virtual photon, there is an explicit zeroth-order contribution

$$C_{L,HH}^{\text{VF},n_f,(0)}(z, Q^2/m_H^2) = e_H^2 z \frac{4m_H^2}{Q^2} \delta\left(z - \frac{Q^2}{Q^2 + m_H^2}\right). \quad (29)$$

This disappears at high Q^2 and the correct zero-mass limit is reached.

Such a zeroth-order coefficient function is implicit in the original ACOT definition of a GM-VFNS, and leads to a peculiar behaviour of $F_L^{c\bar{c}}$ just above $Q^2 = m_c^2$. It is convoluted with the

heavy flavour distribution, which for Q^2 just above m_c^2 is small in magnitude. However, the coefficient function is large near m_c^2 , while the unsubtracted (i.e. FFNS) gluon and singlet-quark coefficient functions are suppressed by a factor of v^3 , where v is the velocity of the heavy quark in the centre-of-mass frame, and are very small for low Q^2 . This means that this zeroth-order heavy-flavour contribution dominates just above $Q^2 = m_c^2$, despite the fact that in a FFNS, where the $c\bar{c}$ pair has to be created, as it must in reality, the contribution is absent.

The contribution from $C_{L,cc}^{\text{VF},n_f,(0)} \otimes (c + \bar{c})$ turns on rapidly just above m_c^2 , dominating other contributions, then dies away as m_H^2/Q^2 becomes small. This leads to a distinct bump in $F_L^{c\bar{c}}(x, Q^2)$ for Q^2 just above m_c^2 , as pointed out in Ref. [86]. In principle this cancels between orders in a properly defined GM-VFNS, as this contribution implicitly appears in the subtraction terms for the gluon and singlet-quark coefficient functions with opposite sign to its explicit contribution. However, the cancellation is imperfect at finite order, and even the partially cancelled contribution dominates at NLO. If this coefficient function is implemented it leads to peculiar behaviour for Q^2 slightly above m_c^2 . At NNLO, where heavy-flavour distributions begin at $Q^2 = m_c^2$ with negative values, the “bump” in $F_L^{c\bar{c}}(x, Q^2)$ is negative, as illustrated in Fig. 18 of Ref. [89], highlighting the unphysical nature of this contribution.

Hence, as in Ref. [86] we choose to ignore the explicit single heavy-quark–photon scattering results. We define the longitudinal sector in what seems to us to be the most physical generalisation of the definition for $F_2^H(x, Q^2)$, as explained in Ref. [72]. The heavy-quark coefficient functions are simply those for the light quarks, with the upper limit of integration moved from 1 to $x_{\text{max}} = Q^2/(Q^2 + 4m_H^2)$. Thus the physical threshold of $W^2 \geq 4m_H^2$ is contained in all terms, and there are no spurious zeroth-order terms. These could only make a contribution if one works in the framework of single heavy quark scattering in the region of low Q^2 where the parton model for the heavy quark is least appropriate. The definition of the SACOT-type scheme [96], and particularly the SACOT(χ) scheme used in the global fits [84], also avoids this undesirable zeroth-order coefficient function. We will discuss our predictions for $F_L(x, Q^2)$ in Section 13.

Charged-current structure functions

The extension of the GM-VFNS to charged currents is most important for the heavy-flavour contribution to neutrino structure function data from CCFR, NuTeV and CHORUS, where the dominant process is $\nu + s \rightarrow \ell^- + c$ (with small Cabibbo mixing) and the charge conjugate process. As such, this is particularly important for an analysis of the dimuon production data from CCFR and NuTeV, discussed in more detail in Section 8. There are also charged-current data from HERA, but this is less precise (though it does not need nuclear corrections, unlike the neutrino data), and is at sufficiently high Q^2 that the heavy quarks are effectively massless, though the GM-VFNS is applied in this case also.

The general procedure for the GM-VFNS for charged-current deep-inelastic scattering works on the same principles as for neutral currents — one can now produce a *single* charm quark from a strange quark so the threshold is now at $x_{\text{max}} = Q^2/(Q^2 + m_c^2)$. However, as explained in Ref. [72], there is a complication because the massive FFNS coefficient functions are not known at $\mathcal{O}(\alpha_S^2)$ (only asymptotic limits [97] have been calculated). These coefficient functions




Figure 4: The gluon coefficient function at $\mathcal{O}(\alpha_S)$ in NC and CC processes.

are needed in our GM-VFNS at low Q^2 at NLO, and at all Q^2 at NNLO — though in the latter case the definition of the GM-VFNS means that the $\ln(Q^2/m_c^2)$ terms are subtracted, and the $\mathcal{O}(m_c^2/Q^2)$ terms die away at high Q^2 , so the GM-VFNS coefficient functions tend to the precisely known massless limits for large Q^2/m_c^2 .

The initial proposal to deal with this, outlined in Ref. [72], was to assume that the mass-dependence in the $\mathcal{O}(\alpha_S^2)$ coefficient functions is the same as for the neutral-current functions, but with the threshold in $4m_c^2$ replaced by a threshold in m_c^2 . It was noted that this meant that the coefficient functions at least satisfy the threshold requirements, and tend smoothly to the correct massless limits, so were very likely to be an improvement on the ZM-VFNS. However, in the course of the analyses performed in the latest global fit we have noticed various complications. One consideration is that the neutrino cross sections are given by expressions of the form

$$\frac{d^2\sigma^{\text{CC}}}{dx dy} \propto \left(1 - y + \frac{y^2}{2} - \frac{M_N xy}{2E_\nu}\right) F_2^{\text{CC}}(x, Q^2) - \frac{y^2}{2} F_L^{\text{CC}}(x, Q^2) \pm y \left(1 - \frac{y}{2}\right) x F_3^{\text{CC}}(x, Q^2), \quad (30)$$

where $y \sim 0.3 - 0.8$. Unlike the general case for neutral currents, where the $F_2(x, Q^2)$ term dominates, all three terms are important for charged currents and in the heavy-flavour contribution there are significant cancellations between them, so all need to be treated very carefully.

Let us first discuss $F_2^{\text{CC}}(x, Q^2)$. At $\mathcal{O}(\alpha_S^2)$ this is dominated by the gluon contribution. The simple prescription, suggested in Ref. [72], gives large corrections to $F_2^{\text{CC}}(x, Q^2)$ at low Q^2 because the lower threshold compared to the neutral-current case leads to a longer convolution length. However, let us consider the comparison of the gluon coefficient functions at $\mathcal{O}(\alpha_S)$, represented by the two diagrams in Fig. 4.

The neutral-current coefficient function is infrared finite and positive. On the other hand the charged-current coefficient function diverges due to the collinear emission of a *light* (i.e. strange) quark. After subtraction of this divergence, via the usual factorisation theorem in the $\overline{\text{MS}}$ scheme, there is an approximate factor of $(1 + 2 \ln(1 - z))$ in the coefficient function at low Q^2 . It is negative for high z even in the FFNS. At higher Q^2 , the finite parts of the neutral-current and charged-current GM-VFNS coefficient functions, after the subtraction of the $\ln(Q^2/m_c^2)$

terms, converge to the same massless expression, but from qualitatively different forms at lower Q^2 .

At $\mathcal{O}(\alpha_S^2)$, the neutral-current FFNS coefficient function is positive with threshold log enhancements. On the other hand, obtaining the approximate charged-current expression by just a change in threshold limits leads to a large contribution, which, bearing in mind the above discussion, is unlikely to be accurate. At $\mathcal{O}(\alpha_S^2)$, the charged-current diagram will still have an extra emission of a strange quark and a corresponding collinear subtraction. So we choose to obtain the charged-current coefficient function by both the change in threshold kinematics and by introducing a $(1 + 2(m_c^2/Q^2) \ln(1 - z))$ factor. This still leads to the correct large Q^2/m_c^2 limit, but is very likely to be a more accurate representation of the low Q^2 region. The same procedure is applied to the similar $\mathcal{O}(\alpha_S^2)$ singlet-quark contribution.

We now consider $xF_3^{\text{CC}}(x, Q^2)$ at $\mathcal{O}(\alpha_S^2)$. For massless quarks this contribution is zero for initial gluons and singlet quarks, but the coefficient function $C_{3,g}^{\text{CC},n_f}(x, m_c^2, Q^2)$ is non-zero for finite Q^2/m_c^2 . However, it must vanish as both $Q^2/m_c^2 \rightarrow \infty$ and $W^2/m_c^2 \rightarrow \infty$. Hence, our model for $C_{3,g}^{\text{CC},n_f,(2)}(x, m_c^2, Q^2)$ is weighted by a factor $m_c^2/\hat{W}^2 = m_c^2 z/((1 - z)Q^2)$, as is the singlet-quark contribution. It is important that a suppression of this type is implemented. Otherwise the contribution is potentially anomalously important at low Q^2 and x .

Finally, there is a complication in the ordering for the longitudinal charged-current heavy flavour production. In the massless limit the lowest-order contribution, simplified by neglecting the Cabibbo-suppressed d -quark contribution, is

$$F_L^{\text{CC},c}(x, Q^2) = 2 \alpha_S (C_{L,g}^{\text{CC}}(x) \otimes g(x, Q^2) + C_{L,q}^{\text{CC}}(x) \otimes s(x, Q^2)). \quad (31)$$

However, for a massive quark there is a zeroth-order contribution

$$F_L^{\text{CC},c}(x, Q^2) = 2 \frac{m_c^2}{m_c^2 + Q^2} \xi s(\xi, Q^2), \quad (32)$$

where $\xi = x(1 + m_c^2/Q^2)$. Note that this is unlike the neutral-current case for $F_L(x, Q^2)$, where there was also a zeroth-order contribution. Here it is due to a real physical process, i.e. $W^+ + s \rightarrow c$, rather than one which only makes sense in the limit where the charm quark is most definitely behaving like a massless parton. Hence, for the charged-current case, the zero-order contribution must be included. This means that there is a difference in orders below and above the transition point, i.e. the FFNS begins at zeroth order whereas the ZM-VFNS begins at first order — opposite to the case for the neutral-current $F_2(x, Q^2)$. Again a choice in ordering must be made. We choose to obtain the correct limits in both regimes along with maintaining continuity. That is, we use (32) to define the LO contribution for $Q^2 < m_c^2$, whereas for $Q^2 > m_c^2$ the LO contribution is defined by

$$F_L^{\text{CC},c}(x, Q^2) = 2 \frac{m_c^2}{m_c^2 + Q^2} \xi s(\xi, Q^2) + 2 \left(1 - \frac{m_c^2}{Q^2}\right) \alpha_S (C_{L,g}(x) \otimes g(x, Q^2) + C_{L,q}(x) \otimes s(x, Q^2)). \quad (33)$$

At high Q^2 the first term naturally dies away leading to the normal massless limit. We can easily generalise the prescription to higher orders by including the next term in the α_S expansion on

both sides of the transition point, with the $(1 - m_c^2/Q^2)$ factor always multiplying the highest-order term in the region above the transition point.

In principle, we could also make some use of the $\mathcal{O}(\alpha_s^3)$ charged-current coefficient functions in the same way as we do for neutral currents. However, the region where they make a significant impact is overwhelmingly in the small x and Q^2 regime accessed only by HERA neutral-current measurements. The amount of modelling required for these terms in charged-current processes is large. Since they are unlikely to have much effect at all, we simply omit them.

With the chosen prescription for charged-current heavy flavour described above, we find that the description of the data, and the resulting parton distributions, are rather stable in going from LO \rightarrow NLO \rightarrow NNLO. This will be discussed further in later sections.

4.4 Intrinsic heavy flavour

Throughout the above discussion we have made the assumption that all heavy quark (c, b) flavour is perturbatively generated from the gluon and lighter flavours. Indeed, this has always been the assumption in our analyses. It is justified up to corrections of $\mathcal{O}(\Lambda_{QCD}^2/m_H^2)$. This potential power-suppressed correction is known as intrinsic heavy flavour. Although it is parametrically small it may be relatively enhanced at high values of x [98], and a plausible measure of the contribution is that at low scales the integrated number density to find intrinsic charm is $\leq 1\%$. There is little within the current global fit which can tightly constrain this possible intrinsic flavour contribution. The possibility of different types of intrinsic charm contribution has been studied recently by CTEQ [99]. They conclude that a global fit can allow up to a factor ~ 3 times more high- x enhanced intrinsic charm than the traditional value of 1%, and that a sea-like contribution can carry a momentum fraction up to 2% of that of the proton (the integrated number density being infinite in this case). Indeed, for the sea-like distribution, they find a preference for a non-zero value, but we believe that this merely compensates for the terms which are systematically absent in the ACOT(χ) scheme at low Q^2 compared to our TR' scheme or to the FFNS. There seems to be no theoretical impetus to consider a sea-like intrinsic heavy flavour to be anything other than tiny. We will consider the possibilities for a high- x enhanced contribution in Section 9. However, in order to do so, we discuss briefly how it fits into the GM-VFNS framework.

Intrinsic flavour and GM-VFNS definitions were discussed in Ref. [100]. We briefly recall them here. Allowing an intrinsic heavy quark distribution actually removes the redundancy in the definition of the coefficient functions in the GM-VFNS, and two different definitions of a GM-VFNS will no longer be identical if formally summed to all orders, though they will only differ by contributions depending on the intrinsic flavour. The reason is as follows. Consider using identical parton distributions, including the intrinsic heavy quarks, in two different flavour schemes. The heavy-quark coefficient functions at each order are different by $\mathcal{O}(m_H^2/Q^2)$. This difference has been constructed to disappear at all orders when combining the parton distributions other than the intrinsic heavy quarks, but will persist for the intrinsic contribution. The intrinsic heavy-flavour distributions are of $\mathcal{O}(\Lambda_{QCD}^2/m_H^2)$, and when combined with the difference in coefficient functions the mass-dependence cancels leading to a difference in structure functions of $\mathcal{O}(\Lambda_{QCD}^2/Q^2)$. It has been shown [101] that for a given GM-VFNS the calculation of

the structure functions is limited in accuracy to $\mathcal{O}(\Lambda_{QCD}^2/Q^2)$. Hence, when including intrinsic charm, the scheme ambiguity is of the same order as the best possible accuracy one can obtain in leading-twist QCD, which is admittedly better than that obtained from ignoring the intrinsic heavy flavour (if it exists) as Q^2 increases above m_H^2 . It is intuitively obvious that better accuracy will be obtained from defining a GM-VFNS where all coefficient functions respect threshold kinematics, e.g. the ACOT(χ) and TR' schemes, than from some older schemes which do not, and which violate physical thresholds when combined with intrinsic heavy-flavour contributions.

5 Global parton analyses

In this section we present a summary of three global PDF analyses performed using the theoretical formalism outlined in Section 3, together with the treatment of heavy flavours given in Section 4. We perform fits at LO, NLO and NNLO starting from input parton distributions parameterised at $Q_0^2 = 1 \text{ GeV}^2$ in the form shown in Eqs. (6)–(12).

5.1 Choice of data sets

The data sets included in the analyses are listed in Table 2, together with the χ^2 values, defined in Section 5.2, corresponding to each individual data set for each of the three fits, and the number of individual data points fitted for each data set. The data sets in Table 2 are ordered according to the type of process. First we have the fixed-target data, which are subdivided into $\ell^\pm N$ structure functions, low-mass Drell–Yan (DY) cross sections, and νN structure functions and dimuon cross sections. Then we list data collected at HERA, and finally data collected at the Tevatron. More detailed discussion of the description of the individual data sets will be given later in Sections 7–12.

For the DIS data we make cuts on the exchanged boson virtuality of $Q^2 \geq 2 \text{ GeV}^2$ and on the squared boson–nucleon centre-of-mass energy, $W^2 = Q^2(1/x - 1) + m_N^2$, of $W^2 \geq 15 \text{ GeV}^2$ at NLO and NNLO, while $W^2 \geq 20 \text{ GeV}^2$ at LO due to the absence of higher-order large- x terms. The cut on W^2 is necessary to exclude data for which potentially large higher-twist contributions are present, such as target mass corrections [115]. For the $\nu(\bar{\nu})N$ xF_3 data we only include data with $W^2 \geq 25 \text{ GeV}^2$ for all three fits. The higher-twist contributions for $F_2^{\nu(\bar{\nu})N}(x, Q^2)$ do not contribute to the Adler sum rule, i.e. $\int_0^1 dx F_2^{\text{HT}}(x, Q^2) = 0$, which means they must be well behaved as $x \rightarrow 0$. However, there is no such restriction on the higher-twist corrections for xF_3 , and renormalon calculations imply that they are large [116, 117]. The CHORUS data extend into the region of expected large higher-twist corrections, so the above cut is necessary to avoid contamination (which is qualitatively similar to that predicted in Refs. [116, 117]).

As in previous MRST fits, data on the deuteron structure function F_2^d , and the ratio F_2^d/F_2^p [103], are corrected for nuclear shadowing effects [118]. No such corrections are applied to the E866/NuSea pd/pp Drell–Yan ratio [108], though we expect such corrections to be small compared to the experimental errors, and it is not clear how such corrections would be applied. For other nuclear target data we apply an improved nuclear correction discussed in detail in Section 7.3. We apply a correction of -3.4% to the luminosity of the published H1 MB 97 data [109] following a luminosity reanalysis [119].

The Tevatron high- p_T jet data are included in the fit at each order. The full NNLO corrections are not known in this case, but the approximation based on threshold corrections [120] are included in the FASTNLO package [75]. There is no guarantee that these give a very good approximation to the full NNLO corrections, but in this case the NLO corrections themselves are of the same order as the systematic uncertainties on the data. The threshold corrections are the only realistic source of large NNLO corrections, so the fact that they provide a correction which is smooth in p_T and small compared to NLO (and compared to the systematic uncertainties on the data) strongly implies that the full NNLO corrections would lead to very little change

Data set	LO	NLO	NNLO
BCDMS $\mu p F_2$ [32]	165 / 153	182 / 163	170 / 163
BCDMS $\mu d F_2$ [102]	162 / 142	190 / 151	188 / 151
NMC $\mu p F_2$ [33]	137 / 115	121 / 123	115 / 123
NMC $\mu d F_2$ [33]	120 / 115	102 / 123	93 / 123
NMC $\mu n/\mu p$ [103]	131 / 137	130 / 148	135 / 148
E665 $\mu p F_2$ [104]	59 / 53	57 / 53	63 / 53
E665 $\mu d F_2$ [104]	49 / 53	53 / 53	63 / 53
SLAC $ep F_2$ [105, 106]	24 / 18	30 / 37	31 / 37
SLAC $ed F_2$ [105, 106]	12 / 18	30 / 38	26 / 38
NMC/BCDMS/SLAC F_L [32–34]	28 / 24	38 / 31	32 / 31
E866/NuSea pp DY [107]	239 / 184	228 / 184	237 / 184
E866/NuSea pd/pp DY [108]	14 / 15	14 / 15	14 / 15
NuTeV $\nu N F_2$ [37]	49 / 49	49 / 53	46 / 53
CHORUS $\nu N F_2$ [38]	21 / 37	26 / 42	29 / 42
NuTeV $\nu N xF_3$ [37]	62 / 45	40 / 45	34 / 45
CHORUS $\nu N xF_3$ [38]	44 / 33	31 / 33	26 / 33
CCFR $\nu N \rightarrow \mu\mu X$ [39]	63 / 86	66 / 86	69 / 86
NuTeV $\nu N \rightarrow \mu\mu X$ [39]	44 / 40	39 / 40	45 / 40
H1 MB 99 e^+p NC [31]	9 / 8	9 / 8	7 / 8
H1 MB 97 e^+p NC [109]	46 / 64	42 / 64	51 / 64
H1 low Q^2 96–97 e^+p NC [109]	54 / 80	44 / 80	45 / 80
H1 high Q^2 98–99 e^-p NC [110]	134 / 126	122 / 126	124 / 126
H1 high Q^2 99–00 e^+p NC [35]	153 / 147	131 / 147	133 / 147
ZEUS SVX 95 e^+p NC [111]	35 / 30	35 / 30	35 / 30
ZEUS 96–97 e^+p NC [112]	118 / 144	86 / 144	86 / 144
ZEUS 98–99 e^-p NC [113]	61 / 92	54 / 92	54 / 92
ZEUS 99–00 e^+p NC [114]	75 / 90	63 / 90	65 / 90
H1 99–00 e^+p CC [35]	28 / 28	29 / 28	29 / 28
ZEUS 99–00 e^+p CC [36]	36 / 30	38 / 30	37 / 30
H1/ZEUS $ep F_2^{\text{charm}}$ [41–47]	110 / 83	107 / 83	95 / 83
H1 99–00 e^+p incl. jets [59]	109 / 24	19 / 24	—
ZEUS 96–97 e^+p incl. jets [57]	88 / 30	30 / 30	—
ZEUS 98–00 $e^\pm p$ incl. jets [58]	102 / 30	17 / 30	—
DØ II $p\bar{p}$ incl. jets [56]	193 / 110	114 / 110	123 / 110
CDF II $p\bar{p}$ incl. jets [54]	143 / 76	56 / 76	54 / 76
CDF II $W \rightarrow \ell\nu$ asym. [48]	50 / 22	29 / 22	30 / 22
DØ II $W \rightarrow \ell\nu$ asym. [49]	23 / 10	25 / 10	25 / 10
DØ II Z rap. [53]	25 / 28	19 / 28	17 / 28
CDF II Z rap. [52]	52 / 29	49 / 29	50 / 29
All data sets	3066 / 2598	2543 / 2699	2480 / 2615

Table 2: The values of $\chi^2/N_{\text{pts.}}$ for the data sets included in the global fits. For the NuTeV $\nu N \rightarrow \mu\mu X$ data, the effective number of degrees of freedom is quoted instead of $N_{\text{pts.}}$ since smearing effects mean that nearby data points are highly correlated [39]. The details of corrections to data, kinematic cuts applied and definitions of χ^2 are contained in the text.

in the PDFs. Since these jet data are the only good direct constraint on the high- x gluon we choose to include them in the NNLO fit judging that the impact of leaving them out would be more detrimental than any inaccuracies in including them without knowing the full NNLO hard cross section. The HERA jet data, for which the NNLO coefficient functions are also unknown at present, are omitted from our NNLO fit. This is because they do not provide as important a constraint as the Tevatron jet data, and also because the NLO corrections are much larger than for Tevatron jet data (due to the lower scales) and there is not even some partial information on the NNLO result available. However, we compare the HERA jet data to results calculated with the NNLO PDFs using the NLO coefficient functions in Section 12, finding surprisingly good agreement.

5.2 χ^2 definition

The global goodness-of-fit measure is defined as $\chi_{\text{global}}^2 = \sum_n \chi_n^2$, where, in all previous MRST fits,

$$\chi_n^2(\{a\}, \mathcal{N}_n) = \sum_{i=1}^{N_{\text{pts}}} \frac{(D_{n,i} - T_{n,i}(\{a\})/\mathcal{N}_n)^2}{(\sigma_{n,i}^{\text{uncorr.}})^2 + \sum_k (\sigma_{n,k,i}^{\text{corr.}})^2} + \chi_{\mathcal{N}_n}^2. \quad (34)$$

Here, n labels a particular data set, or a combination of data sets, with a common (fitted) normalisation \mathcal{N}_n , i labels the number of individual data points in that data set, and k labels the individual correlated systematic errors for a particular data set. The individual data points $D_{n,i}$ have uncorrelated (statistical and systematic) errors $\sigma_{n,i}^{\text{uncorr.}}$ and correlated systematic errors $\sigma_{n,k,i}^{\text{corr.}}$. The theory prediction $T_{n,i}(\{a\})$ depends on the input PDF parameters $\{a\}$. The MRST fits added all uncorrelated and correlated experimental errors in quadrature in the χ^2 definition (34), except for Tevatron jet production. Here, the full correlated error information was used but the procedure was more complicated; see Section 12. The second term in (34) is a possible penalty term to account for deviations of the fitted data set normalisations \mathcal{N}_n from the nominal values of 1. (In practice, the MRST fits took $\chi_{\mathcal{N}_n}^2 = 0$, although constraints were applied as explained below.) We will now discuss improvements to the treatment of data set normalisations made in the present analysis, then we will discuss improvements in the treatment of correlated systematic errors.

Treatment of data set normalisations

In the MRST analyses, the normalisations of several data sets were floated and then fixed, checking that these normalisations lay inside the quoted one-sigma error range. For example, the H1 data sets from different running periods were all taken to have a common normalisation of exactly 1, while the ZEUS data sets from different running periods were all taken to have a common normalisation of 0.98. In general, HERA data from different runs are allowed to have different normalisations, therefore we now split the H1 and ZEUS data into different running periods, and allow a different normalisation for each of these running periods.¹⁰ The normalisa-

¹⁰We take the normalisations of the H1 and ZEUS data on $F_2^{e\bar{e}}$ to be fixed at 1, although here the experimental errors are anyway much larger than for the total reduced cross sections, and the $F_2^{e\bar{e}}$ data do not provide a strong

tions of the data sets are now all allowed to go free at the same time as the PDF parameters, including in the determination of the PDF uncertainties discussed in Section 6.

Note that the (fitted) normalisation parameters \mathcal{N}_n are included in (34) in the correct manner to avoid bias in the result, i.e. if the central value of the data is rescaled by a factor \mathcal{N}_n then so are the uncertainties. If fitting a single data set with no normalisation constraint on the theory we would automatically obtain $\mathcal{N}_n = 1$. However, the fit does allow data sets to have different relative normalisations, and additionally the sum rules on the PDFs do influence the normalisation to some extent.

The MRST analyses did not include any penalty term in the χ^2 definition for data set normalisations \mathcal{N}_n differing from the nominal values of 1. The usual choice for the penalty term for the data set normalisations is

$$\chi_{\mathcal{N}_n}^2 = \left(\frac{1 - \mathcal{N}_n}{\sigma_n^{\mathcal{N}}} \right)^2, \quad (35)$$

where $\sigma_n^{\mathcal{N}}$ is the one-sigma normalisation error for data set n . Alternatively, it has been proposed that normalisation uncertainties should behave according to a log-normal distribution, where the penalty term is [121]

$$\chi_{\mathcal{N}_n}^2 = \ln(\mathcal{N}_n) \left(3 + \frac{\ln(\mathcal{N}_n)}{\ln(1 + (\sigma_n^{\mathcal{N}})^2)} \right), \quad (36)$$

which reduces to the usual quadratic term (35) for small $\sigma_n^{\mathcal{N}}$ and for small deviations of \mathcal{N}_n from 1. We find that, using (35), the best-fit data set normalisations tend to stray outside their nominal one-sigma range, with all the largest shifts being in the downwards direction. It has long been known that both LO and NLO fits would prefer to have more than 100% momentum for the PDFs if allowed (e.g. at NLO in Ref. [122], the extra momentum in the fit regions for the conservative NLO PDFs in Ref. [17], and the large momentum violation at LO seen in Ref. [66]) so imposing momentum conservation on PDFs conversely leads to a preference for lower data set normalisations. However, this is a diminishing effect as we increase the order of the QCD calculation, implying that it is an artifact of an incomplete theory at lower orders, and as such is an undesirable systematic effect. Additionally, it has been claimed that normalisation uncertainties are expected to behave more like a box shape than the usual Gaussian behaviour [123], and indeed, the term in (36) does alter the shape of the penalty in the downwards direction in this type of manner, albeit to a very small degree for errors of only a few percent. Hence, taking into account the systematic effect at low perturbative orders and that normalisations are very unlikely to be perfectly Gaussian, we use a more severe *quartic* penalty term for the normalisations, i.e.

$$\chi_{\mathcal{N}_n}^2 = \left(\frac{1 - \mathcal{N}_n}{\sigma_n^{\mathcal{N}}} \right)^4. \quad (37)$$

This binds the normalisations more strongly to the range $\mathcal{N}_n \in [1 - \sigma_n^{\mathcal{N}}, 1 + \sigma_n^{\mathcal{N}}]$, but in practice it is far more the case that it stops them from floating too low. If more than one data set has a common normalisation, the penalty term $\chi_{\mathcal{N}_n}^2$ is divided amongst those data sets according to the number of data points.

constraint on the PDFs.

The resulting (re)normalisation factors \mathcal{N}_n for each fit are listed in Table 3.

Correlated systematic errors

The insensitivity of both the central values of the PDFs and their uncertainties to the complete inclusion of correlation information for the DIS data was confirmed in the benchmark PDF fits in the proceedings of the HERA–LHC workshop [71] and an extensive discussion of the effect of correlations for HERA data was presented in the appendix of Ref. [14]. To summarise the latter, the main effect was by far the systematic shifts due to changes in normalisation which we now treat in full (indeed a strong hint of the necessity for the normalisation correction [119] of the H1 MB 97 data [109] was seen). Beyond this, the introduction of full correlations led to rather small changes and those were, in general, precisely such as to flatten the evolution of $F_2(x, Q^2)$ allowing the gluon to reduce. As previous comments on normalisations and momentum conservation suggest, this is itself likely correlated to shortcomings of the theory at low orders. Hence, we maintain that the lack of the correlations has little effect on our results, and continue to make this simplification. This is particularly the case since we note that for the preliminary averaged HERA cross section data [124], sources of correlated uncertainty are often dramatically reduced (sometimes by a factor of 3–4), and in the averaging between H1 and ZEUS the data points frequently move relative to each other in a rather different manner than a data set does relative to theory in a fit. We intend to include full correlations for the averaged HERA data, when they are published and will be better understood, but at this stage they will also be very small indeed. We also note that in some cases a textbook treatment of correlated uncertainties can lead to peculiar results. Indeed in Ref. [125] it was seen that the high- y turnover in reduced cross section data due to the influence of $F_L(x, Q^2)$ could be eliminated completely in a fit by letting the correlated systematic uncertainty due to the photoproduction background move by about two-sigma. This seems unlikely (though not impossible), and it is certainly the case that the distribution of uncertainties in this background are far from Gaussian, so the conventional treatment may give surprising results. It is definitely the case that the absolute χ^2 values for the best fit do change depending on how the correlations are treated, but given that our new tolerance determination, described in detail in Section 6.2, relies on changes in χ^2 relative to the best fit, we are confident that we are even less sensitive to such details than in previous studies, e.g. the aforementioned [16, 71].

Nevertheless, for selected newly-added data sets we do include the full correlated error information. Instead of (34), the χ^2 is given by [126]

$$\chi_n^2(\{a\}, \mathcal{N}_n) = \sum_{i=1}^{N_{\text{pts}}} \left(\frac{\hat{D}_{n,i} - T_{n,i}(\{a\})/\mathcal{N}_n}{\sigma_{n,i}^{\text{uncorr.}}} \right)^2 + \sum_{k=1}^{N_{\text{corr.}}} r_{n,k}^2 + \chi_{\mathcal{N}_n}^2, \quad (38)$$

where $\hat{D}_{n,i} \equiv D_{n,i} - \sum_{k=1}^{N_{\text{corr.}}} r_{n,k} \sigma_{n,k,i}^{\text{corr.}}$ are the data points allowed to shift by the systematic errors in order to give the best fit. Minimising χ_n^2 with respect to $r_{n,k}$ gives the analytic result that [126]

$$r_{n,k}(\{a\}, \mathcal{N}_n) = \sum_{k'=1}^{N_{\text{corr.}}} (A^{-1})_{kk'} B_{k'}(\{a\}, \mathcal{N}_n), \quad (39)$$

Data set	$\sigma_n^{\mathcal{N}}$	LO	NLO	NNLO
BCDMS μp F_2 [32]	3%	0.9667	0.9644	0.9678
BCDMS μd F_2 [102]	3%	0.9667	0.9644	0.9678
NMC μp F_2 [33]	2%	1.0083	0.9982	0.9999
NMC μd F_2 [33]	2%	1.0083	0.9982	0.9999
NMC $\mu n/\mu p$ [103]	—	1	1	1
E665 μp F_2 [104]	1.85%	1.0146	1.0052	1.0024
E665 μd F_2 [104]	1.85%	1.0146	1.0052	1.0024
SLAC ep F_2 [105, 106]	1.9%	1.0227	1.0125	1.0078
SLAC ed F_2 [105, 106]	1.9%	1.0227	1.0125	1.0078
NMC/BCDMS/SLAC F_L [32–34]	—	1	1	1
E866/NuSea pp DY [107]	6.5%	1.0629	1.0086	1.0868
E866/NuSea pd/pp DY [108]	—	1	1	1
NuTeV νN F_2 [37]	2.1%	0.9987	0.9997	0.9992
CHORUS νN F_2 [38]	2.1%	0.9987	0.9997	0.9992
NuTeV νN xF_3 [37]	2.1%	0.9987	0.9997	0.9992
CHORUS νN xF_3 [38]	2.1%	0.9987	0.9997	0.9992
CCFR $\nu N \rightarrow \mu\mu X$ [39]	2.1%	0.9987	0.9997	0.9992
NuTeV $\nu N \rightarrow \mu\mu X$ [39]	2.1%	0.9987	0.9997	0.9992
H1 MB 99 e^+p NC [31]	1.3%	0.9861	1.0098	1.0090
H1 MB 97 e^+p NC [109]	1.5%	0.9863	0.9921	0.9953
H1 low Q^2 96–97 e^+p NC [109]	1.7%	1.0029	1.0095	1.0172
H1 high Q^2 98–99 e^-p NC [110]	1.8%	0.9782	0.9851	0.9860
H1 high Q^2 99–00 e^+p NC [35]	1.5%	0.9762	0.9834	0.9842
ZEUS SVX 95 e^+p NC [111]	1.5%	0.9944	0.9948	1.0004
ZEUS 96–97 e^+p NC [112]	2%	0.9735	0.9811	0.9871
ZEUS 98–99 e^-p NC [113]	1.8%	0.9771	0.9855	0.9862
ZEUS 99–00 e^+p NC [114]	2.5%	0.9656	0.9761	0.9762
H1 99–00 e^+p CC [35]	1.5%	0.9762	0.9834	0.9842
ZEUS 99–00 e^+p CC [36]	2.5%	0.9656	0.9761	0.9762
H1/ZEUS ep F_2^{charm} [41–47]	—	1	1	1
H1 99–00 e^+p incl. jets [59]	1.5%	0.9762	0.9834	—
ZEUS 96–97 e^+p incl. jets [57]	2%	0.9735	0.9811	—
ZEUS 98–00 $e^\pm p$ incl. jets [58]	2.5%	0.9656	0.9761	—
DØ II $p\bar{p}$ incl. jets [56]	6.1%	0.9353	1.0596	1.0759
CDF II $p\bar{p}$ incl. jets [54]	5.8%	0.8779	0.9646	0.9900
CDF II $W \rightarrow \ell\nu$ asym. [48]	—	1	1	1
DØ II $W \rightarrow \ell\nu$ asym. [49]	—	1	1	1
DØ II Z rap. [53]	—	1	1	1
CDF II Z rap. [52]	5.8%	0.8779	0.9646	0.9900

Table 3: The fitted normalisations \mathcal{N}_n of the data sets included in the global fit, together with the one-sigma normalisation errors, $\sigma_n^{\mathcal{N}}$, for each data set n .

where

$$A_{kk'} = \delta_{kk'} + \sum_{i=1}^{N_{\text{pts.}}} \frac{\sigma_{n,k,i}^{\text{corr.}} \sigma_{n,k',i}^{\text{corr.}}}{\sigma_{n,i}^{\text{uncorr.}}}, \quad B_k(\{a\}, \mathcal{N}_n) = \sum_{i=1}^{N_{\text{pts.}}} \frac{\sigma_{n,k,i}^{\text{corr.}} (D_{n,i} - T_{n,i}(\{a\})/\mathcal{N}_n)}{(\sigma_{n,i}^{\text{uncorr.}})^2}. \quad (40)$$

Therefore, the optimal shifts of the data points by the systematic errors are solved for analytically, while the input PDF parameters $\{a\}$, together with the data normalisations \mathcal{N}_n , must be determined by numerical minimisation of χ_{global}^2 in the usual way.

In practice, we use (38) only for the jet-energy-scale uncertainties in the ZEUS inclusive jet data [57, 58], the 4 contributions to the correlated uncertainty in the H1 inclusive jet data [59], the 16 sources of correlated uncertainty in the CDF Run II inclusive jet data [54], the 22 sources of correlated uncertainty in the DØ Run II inclusive jet data [56], and the 6 sources of correlated uncertainty in the CDF Run II Z rapidity distribution [52]. In many other cases the uncertainties (other than normalisation) are presented without any information on correlations between systematic uncertainties, or are assumed to be uncorrelated. For the remainder of the data sets, the systematic errors (other than the normalisation error) are simply treated as being uncorrelated and added in quadrature with the statistical uncertainties, i.e. we use (34) with the penalty term (37).

Minimisation of χ^2 and calculation of Hessian matrix

To determine the best fit at NLO and NNLO we need to minimise the χ_{global}^2 with respect to 28 free input PDF parameters, together with $\alpha_S(Q_0^2)$, 3 parameters (r_1, r_2, r_3) associated with nuclear corrections (see Section 7), and 17 different data set normalisations, giving a total of 49 free parameters. This is a difficult task and is unlikely to be possible with the widely-used MINUIT package [127], where the practical maximum number of free parameters is around 15. Indeed, the CTEQ group have found it necessary to extend MINUIT to use an improved iterative method in the numerical calculation of the Hessian matrix [128]. Other fitting groups, such as Alekhin and H1/ZEUS, use far fewer PDF parameters, and so presumably do not suffer from the same problems when using the standard MINUIT package. As in previous MRS/MRST analyses [2–21], we instead use the Levenberg–Marquardt method [129, 130] described in Numerical Recipes [131], which combines the advantages of the inverse-Hessian method and the steepest descent method for minimisation. The method requires knowledge of the gradient and Hessian matrix of the χ_{global}^2 , which we provide partially as analytic expressions in our fitting code, only using numerical finite-difference computations for the derivatives of the theory predictions with respect to the fitted parameters. A linearisation approximation is made in the calculation of the Hessian matrix to avoid the presence of potentially destabilising second-derivative terms [131]. For example, the contribution to the global Hessian matrix $H_{lm} = \sum_n H_{lm}^n$ from a data set n with χ_n^2 defined by (34) is

$$H_{lm}^n \equiv \frac{1}{2} \frac{\partial \chi_n^2}{\partial a_l \partial a_m} = \sum_{i=1}^{N_{\text{pts.}}} \frac{1}{(\sigma_{n,i}^{\text{uncorr.}})^2 + \sum_k (\sigma_{n,k,i}^{\text{corr.}})^2} \frac{\partial T_{n,i}(\{a\})/\mathcal{N}_n}{\partial a_l} \frac{\partial T_{n,i}(\{a\})/\mathcal{N}_n}{\partial a_m}, \quad (41)$$

while in the more complicated case of treating correlated systematic errors (38), the corresponding expression is

$$H_{lm}^n = \sum_{i=1}^{N_{\text{pts.}}} \frac{1}{(\sigma_{n,i}^{\text{uncorr.}})^2} \frac{\partial \hat{T}_{n,i}}{\partial a_l} \frac{\partial \hat{T}_{n,i}}{\partial a_m} + \sum_{k=1}^{N_{\text{corr.}}} \frac{\partial r_{n,k}}{\partial a_l} \frac{\partial r_{n,k}}{\partial a_m}, \quad (42)$$

where $\hat{T}_{n,i}(\{a\}, \mathcal{N}_n) \equiv T_{n,i}(\{a\})/\mathcal{N}_n + \sum_{k=1}^{N_{\text{corr.}}} r_{n,k}(\{a\}, \mathcal{N}_n) \sigma_{n,k,i}^{\text{corr.}}$. Similar expressions are used for the elements of the Hessian matrix corresponding to data set normalisations \mathcal{N}_n .

5.3 Discussion of fit results

For each of the three fits, the optimal values of the input PDF parameters of Eqs. (6)–(12) and of the QCD coupling α_S are given in Table 4.

Notice that the LO fit gives a significantly larger $\alpha_S(M_Z^2)$ than the world-average value of 0.1176 ± 0.0020 [132], and that the parametrisation of the input gluon is simpler since the fit never shows any tendency for the gluon to go negative, or even to turn over at small x . Moreover, the LO fit gives a much worse description of the data than that obtained at NLO and NNLO. It fails for HERA structure functions because the evolution is too slow on account of missing higher-order small- x enhanced terms in the quark–gluon splitting function. It fails for some fixed-target data due to the absence of higher-order large- x enhanced terms in the coefficient functions. For the HERA jets it is missing large NLO corrections to the cross sections. For Tevatron jets we find that the LO evolution cannot give, in detail, the correct shape to fit the data, contrary to previous findings for the less precise Run I data which were fit very well [15]. It is also missing large NLO corrections to the cross section for the Drell–Yan data, and as before we have applied a K -factor (previously fixed to 1.3), but have improved this to $(1 + \alpha_S(M^2) C_F \pi/2)$, which improves the fit quality considerably. The deficiencies due to higher-order terms cannot be completely mimicked by the very large coupling α_S (although this works quite well at high x) and a large small- x gluon, even when the normalisations of many data sets choose to lie rather low. So, in conclusion, the LO fit is unsatisfactory. On the other hand, if required, it is possible to obtain a modified LO PDF set which gives a much improved description of the data by relaxing the momentum conservation constraint and using the NLO coupling α_S [66], and this can be even further improved by a modification of the scale of the coupling [67].

The quality of the fit at NLO is much better, with no obviously badly fit data sets (except to some extent the DØ $W \rightarrow \ell \nu$ charge asymmetry, which will be discussed later). Again the normalisations tend to lie a little low, though there are systematic differences between data sets, e.g. the SLAC DIS data are normalised about 1.5% higher than NMC, which is 3% higher than BCDMS, and in general H1 sets are 1% or so higher in normalisation than ZEUS. These trends have long been known. We also see that the fitted normalisation of the DØ Run II jet data is about 8% higher than that for CDF Run II data (in the latter case the normalisation of the inclusive jet data is tied to the normalisation of the Z rapidity distribution data), but the luminosity uncertainties in each are compatible with this. The fit quality at NNLO is almost identical to that at NLO, and is rather similar for most data sets. However, it is striking that the normalisations are clustered about unity much more symmetrically in this case, with those

Parameter	LO	NLO	NNLO
$\alpha_S(Q_0^2)$	0.68183	0.49128	0.45077
$\alpha_S(M_Z^2)$	0.13939	0.12018	0.11707
A_u	1.4335	0.25871	0.22250
η_1	0.45232 $^{+0.022}_{-0.018}$	0.29065 $^{+0.019}_{-0.013}$	0.27871 $^{+0.018}_{-0.014}$
η_2	3.0409 $^{+0.079}_{-0.067}$	3.2432 $^{+0.062}_{-0.039}$	3.3627 $^{+0.061}_{-0.044}$
ϵ_u	-2.3737 $^{+0.54}_{-0.48}$	4.0603 $^{+1.6}_{-2.3}$	4.4343 $^{+2.4}_{-2.7}$
γ_u	8.9924	30.687	38.599
A_d	5.0903	12.288	17.938
η_3	0.71978 $^{+0.057}_{-0.082}$	0.96809 $^{+0.11}_{-0.11}$	1.0839 $^{+0.12}_{-0.11}$
$\eta_4 - \eta_2$	2.0835 $^{+0.32}_{-0.45}$	2.7003 $^{+0.50}_{-0.52}$	2.7865 $^{+0.50}_{-0.44}$
ϵ_d	-4.3654 $^{+0.28}_{-0.22}$	-3.8911 $^{+0.31}_{-0.29}$	-3.6387 $^{+0.27}_{-0.28}$
γ_d	7.4730	6.0542	5.2577
A_S	0.59964 $^{+0.036}_{-0.030}$	0.31620 $^{+0.030}_{-0.021}$	0.64942 $^{+0.047}_{-0.041}$
δ_S	-0.16276	-0.21515	-0.11912
η_S	8.8801 $^{+0.33}_{-0.33}$	9.2726 $^{+0.23}_{-0.33}$	9.4189 $^{+0.25}_{-0.33}$
ϵ_S	-2.9012 $^{+0.33}_{-0.37}$	-2.6022 $^{+0.71}_{-0.96}$	-2.6287 $^{+0.49}_{-0.51}$
γ_S	16.865	30.785	18.065
$\int_0^1 dx \Delta(x, Q_0^2)$	0.091031 $^{+0.012}_{-0.009}$	0.087673 $^{+0.013}_{-0.011}$	0.078167 $^{+0.012}_{-0.0091}$
A_Δ	8.9413	8.1084	16.244
η_Δ	1.8760 $^{+0.24}_{-0.30}$	1.8691 $^{+0.23}_{-0.32}$	2.0741 $^{+0.18}_{-0.35}$
γ_Δ	8.4703 $^{+2.0}_{-0.3}$	13.609 $^{+1.1}_{-0.6}$	6.7640 $^{+0.77}_{-0.41}$
δ_Δ	-36.507	-59.289	-36.090
A_g	0.0012216	1.0805	3.4055
δ_g	-0.83657 $^{+0.15}_{-0.14}$	-0.42848 $^{+0.066}_{-0.057}$	-0.12178 $^{+0.23}_{-0.16}$
η_g	2.3882 $^{+0.51}_{-0.50}$	3.0225 $^{+0.43}_{-0.36}$	2.9278 $^{+0.68}_{-0.41}$
ϵ_g	-38.997 $^{+36}_{-35}$	-2.2922	-2.3210
γ_g	1445.5 $^{+880}_{-750}$	3.4894	1.9233
$A_{g'}$	—	-1.1168	-1.6189
$\delta_{g'}$	—	-0.42776 $^{+0.053}_{-0.047}$	-0.23999 $^{+0.14}_{-0.10}$
$\eta_{g'}$	—	32.869 $^{+6.5}_{-5.9}$	24.792 $^{+6.5}_{-5.2}$
A_+	0.10302 $^{+0.029}_{-0.017}$	0.047915 $^{+0.0095}_{-0.0076}$	0.10455 $^{+0.019}_{-0.016}$
η_+	13.242 $^{+2.9}_{-1.4}$	9.7466 $^{+1.0}_{-0.8}$	9.8689 $^{+1.0}_{-0.6}$
A_-	-0.011523 $^{+0.009}_{-0.018}$	-0.011629 $^{+0.009}_{-0.023}$	-0.0093692 $^{+0.006}_{-0.024}$
η_-	10.285 $^{+16}_{-6}$	11.261 $^{+22}_{-6}$	9.5783 $^{+26}_{-5}$
x_0	0.017414	0.016050	0.018556
r_1	-0.39484	-0.57631	-0.80834
r_2	-1.0719	0.81878	1.2669
r_3	-0.28973	-0.083208	0.15098

Table 4: The optimal values of α_S and the input PDF parameters at $Q_0^2 = 1 \text{ GeV}^2$ determined from the global analysis. The one-sigma errors are calculated using (51) and (52) using the 68% C.L. tolerance discussed in Section 6, and are shown only for the 20 parameters allowed to go free when determining the eigenvector PDF sets. The parameters A_u , A_d , A_g and x_0 are determined from sum rules and are not fitted parameters. Similarly, A_Δ is determined from $\int_0^1 dx \Delta(x, Q_0^2)$. The three parameters r_i , defined in (73), are associated with the nuclear corrections to the neutrino data; see Section 7.3. The parameter values are given to five significant figures solely for accuracy in the case of reproduction of the PDFs.

for the E866/NuSea Drell–Yan cross sections and the DØ Run II jet data more than one-sigma above unity. At NLO and NNLO the values of $\alpha_S(M_Z^2)$ of 0.1202 and 0.1171 are both consistent with the world average value of 0.1176 ± 0.0020 [132], though the NLO value is just a little high, suggesting a partial compensation for missing higher orders. A detailed study of the uncertainty in α_S awaits a further publication, but previous estimates of an uncertainty on our determination of $\alpha_S(M_Z^2)$ of ± 0.002 from data and ± 0.003 from theory are unlikely to change very significantly.

The errors that are shown on the input parameters in Table 4 are calculated with (51) and (52) using the 68% C.L. tolerance discussed in Section 6, and are given only for the 20 parameters allowed to go free when determining the eigenvector PDF sets. It is important to note that although they give some idea of PDF uncertainties they are only quantitatively meaningful when the information on the correlations between parameters is included, and this is most easily presented and utilised in terms of orthogonal eigenvectors. The eigenvector PDF sets are defined and discussed in Section 6.1. Indeed, in Section 6 we describe, in detail, the procedure that we use to determine the uncertainty on the parton distributions and on the predictions of cross sections. We then turn to the more detailed description of the various data sets that are obtained in the global fits in Sections 7–12.

6 Error propagation from experimental data points

Uncertainties in global PDF analyses can be divided into two general categories, which may loosely be called “theoretical” and “experimental”. The *theoretical* errors include those associated with the choice of the form of the input parameterisation, the neglected higher-order QCD (including enhanced $\ln(1/x)$ or $\ln(1-x)$ terms) and electroweak corrections, parton recombination and other higher-twist corrections, the choice of data sets and kinematic cuts, the choice of nuclear corrections for the neutrino-initiated data and the treatment of heavy flavours. These uncertainties are often difficult to quantify *a priori* until a better calculation or prescription becomes available, but see the attempt made in Ref. [17].

On the other hand, in principle there are well-defined procedures for propagating *experimental* uncertainties on the fitted data points through to the PDF uncertainties. Three main methods have been used to do this.

1. The *Lagrange multiplier* method [128, 133], which does not rely on linear error propagation, but requires the ability to perform a global fit.
2. The *Hessian* method [134], which is based on linear error propagation and involves the production of eigenvector PDF sets suitable for convenient use by the end user.
3. The use of *Monte Carlo sampling* [135, 136], which has recently been used in conjunction with neural networks to determine NLO PDFs from a DIS-only fit [137].

The first two methods were originally used by CTEQ [128, 133, 134] and then by MRST [16]. In principle, the Lagrange multiplier method is superior to the Hessian approach, but it suffers from the enormous practical disadvantage that a series of new global fits has to be done every time one considers a new quantity. Fortunately, it turns out that for those quantities that have been considered by both methods, the uncertainties have been found to be comparable. We therefore concentrate here on the most commonly used Hessian method. We will not discuss the use of Monte Carlo sampling in this paper.

6.1 Review of Hessian method in global PDF analyses

The basic procedure is discussed in detail in Refs. [16, 134]. Here, we briefly review the important points. The traditional propagation of experimental uncertainties assumes that the global goodness-of-fit quantity, χ_{global}^2 , is quadratic about the global minimum, which has parameters $\{a_1^0, \dots, a_n^0\}$. In this case we can write

$$\Delta\chi_{\text{global}}^2 \equiv \chi_{\text{global}}^2 - \chi_{\text{min}}^2 = \sum_{i,j=1}^n H_{ij}(a_i - a_i^0)(a_j - a_j^0), \quad (43)$$

where the Hessian matrix H has components

$$H_{ij} = \frac{1}{2} \frac{\partial^2 \chi_{\text{global}}^2}{\partial a_i \partial a_j} \bigg|_{\text{min}}. \quad (44)$$

The uncertainty on a quantity $F(\{a_i\})$ is then obtained from linear error propagation:

$$\Delta F = T \sqrt{\sum_{i,j=1}^n \frac{\partial F}{\partial a_i} C_{ij} \frac{\partial F}{\partial a_j}}, \quad (45)$$

where $C \equiv H^{-1}$ is the covariance matrix, also known as the error matrix, and $T = (\Delta\chi_{\text{global}}^2)^{1/2}$ is the tolerance for the required confidence interval. This formula (45) has the disadvantage that PDF uncertainties are not readily calculable for general observables, since the derivative of the observable F with respect to each parameter a_i is needed.

It is convenient to diagonalise the covariance (or Hessian) matrix [16, 134], and work in terms of the eigenvectors and eigenvalues. Since the covariance matrix is symmetric it has a set of *orthonormal* eigenvectors \mathbf{v}_k defined by

$$\sum_{j=1}^n C_{ij} v_{jk} = \lambda_k v_{ik}, \quad (46)$$

where λ_k is the k th eigenvalue and v_{ik} is the i th component of the k th orthonormal eigenvector ($k = 1, \dots, n$). The parameter displacements from the global minimum can be expanded in a basis of rescaled eigenvectors $e_{ik} \equiv \sqrt{\lambda_k} v_{ik}$, that is,

$$a_i - a_i^0 = \sum_{k=1}^n e_{ik} z_k. \quad (47)$$

Then it can be shown, using the orthonormality of \mathbf{v}_k , that (43) reduces to

$$\chi_{\text{global}}^2 = \chi_{\text{min}}^2 + \sum_{k=1}^n z_k^2, \quad (48)$$

that is, $\sum_{k=1}^n z_k^2 \leq T^2$ is the interior of a hypersphere of radius T . Pairs of eigenvector PDF sets S_k^\pm can then be produced to span this hypersphere, at a fixed value of α_S , with parameters given by

$$a_i(S_k^\pm) = a_i^0 \pm t e_{ik}, \quad (49)$$

with t adjusted to give the desired $T = (\Delta\chi_{\text{global}}^2)^{1/2}$. In the quadratic approximation, $t = T$. For the larger eigenvalues λ_k , where there are significant deviations from the ideal quadratic behaviour, t is adjusted iteratively to give the desired value of T . Then uncertainties on a quantity F , which may be an individual PDF at particular values of x and Q^2 , or a derived quantity such as a cross section, can be calculated with¹¹

$$\Delta F = \frac{1}{2} \sqrt{\sum_{k=1}^n [F(S_k^+) - F(S_k^-)]^2}, \quad (50)$$

¹¹It can be shown that (45) reduces to (50) in the quadratic approximation ($t = T$) [16, 134], but we treat (50) as the fundamental definition in the departure of this ideal limit. In this paper, we will generally use (51) and (52) to calculate asymmetric PDF uncertainties.

or asymmetric errors can be calculated with

$$(\Delta F)_+ = \sqrt{\sum_{k=1}^n \left\{ \max \left[F(S_k^+) - F(S_0), F(S_k^-) - F(S_0), 0 \right] \right\}^2}, \quad (51)$$

$$(\Delta F)_- = \sqrt{\sum_{k=1}^n \left\{ \max \left[F(S_0) - F(S_k^+), F(S_0) - F(S_k^-), 0 \right] \right\}^2}, \quad (52)$$

where S_0 is the central PDF set. Correlations between two quantities can also be calculated; see, for example, Ref. [87]. Defining a correlation cosine between two quantities F and G ,

$$\cos \phi_{FG} = \frac{1}{4 \Delta F \Delta G} \sum_{k=1}^n \left[F(S_k^+) - F(S_k^-) \right] \left[G(S_k^+) - G(S_k^-) \right], \quad (53)$$

where the uncertainties ΔF and ΔG are calculated using (50), then values of $\cos \phi_{FG} \approx 1$ mean that F and G are correlated, values of ≈ -1 mean that they are anticorrelated, while values of ≈ 0 mean that they are uncorrelated. A tolerance ellipse in the F - G plane can then be defined by the two parametric equations:

$$F = F(S_0) + \Delta F \cos \theta, \quad (54)$$

$$G = G(S_0) + \Delta G \cos(\theta + \phi_{FG}), \quad (55)$$

where $\theta \in [0, 2\pi]$. We will show examples in Section 15. Note that using e.g. 90% C.L. PDFs results in a probability less than 90% for the 2-D tolerance ellipse [87, 131, 132].

To determine the “best-fit” parameters we allow all the input PDF parameters of Eqs. (6)–(12) to vary. However, when investigating in detail the small departures from the global minimum we notice a certain amount of redundancy in parameters. A striking example is for the NLO and NNLO parameterisations of the gluon distribution (10). Small changes in the values of three of the parameters can be compensated almost exactly by changes in the remaining four. This high degree of correlation between parameters means that very small changes in χ^2 will be obtained. However, at some point the compensation starts to fail significantly and the χ^2 increases dramatically. Hence, the redundancy leads to a severe breaking of the quadratic behaviour in $\Delta\chi^2$, and some very flat directions in the eigenvector space (that is, very large eigenvalues of the covariance matrix) and cubic, quartic etc. terms dominate. During the process of diagonalisation this bad behaviour feeds through into the whole set of eigenvectors to some extent. Therefore, in order that the Hessian method works at all well we have to lessen the redundancy in the input parameters. In order to do this we simply fix some of the parameters at their best-fit values, so that the Hessian matrix only depends on a subset of parameters that are sufficiently independent that the quadratic approximation is reasonable. We finish up with 20 reasonably well-behaved eigenvectors in total, i.e. those coming from the combinations of the 20 parameters that are assigned errors in Table 4. However, we emphasise that the other parameters are fixed at their best-fit values, rather than simply set to zero. The problem of a

MSTW 2008 NLO PDF fit

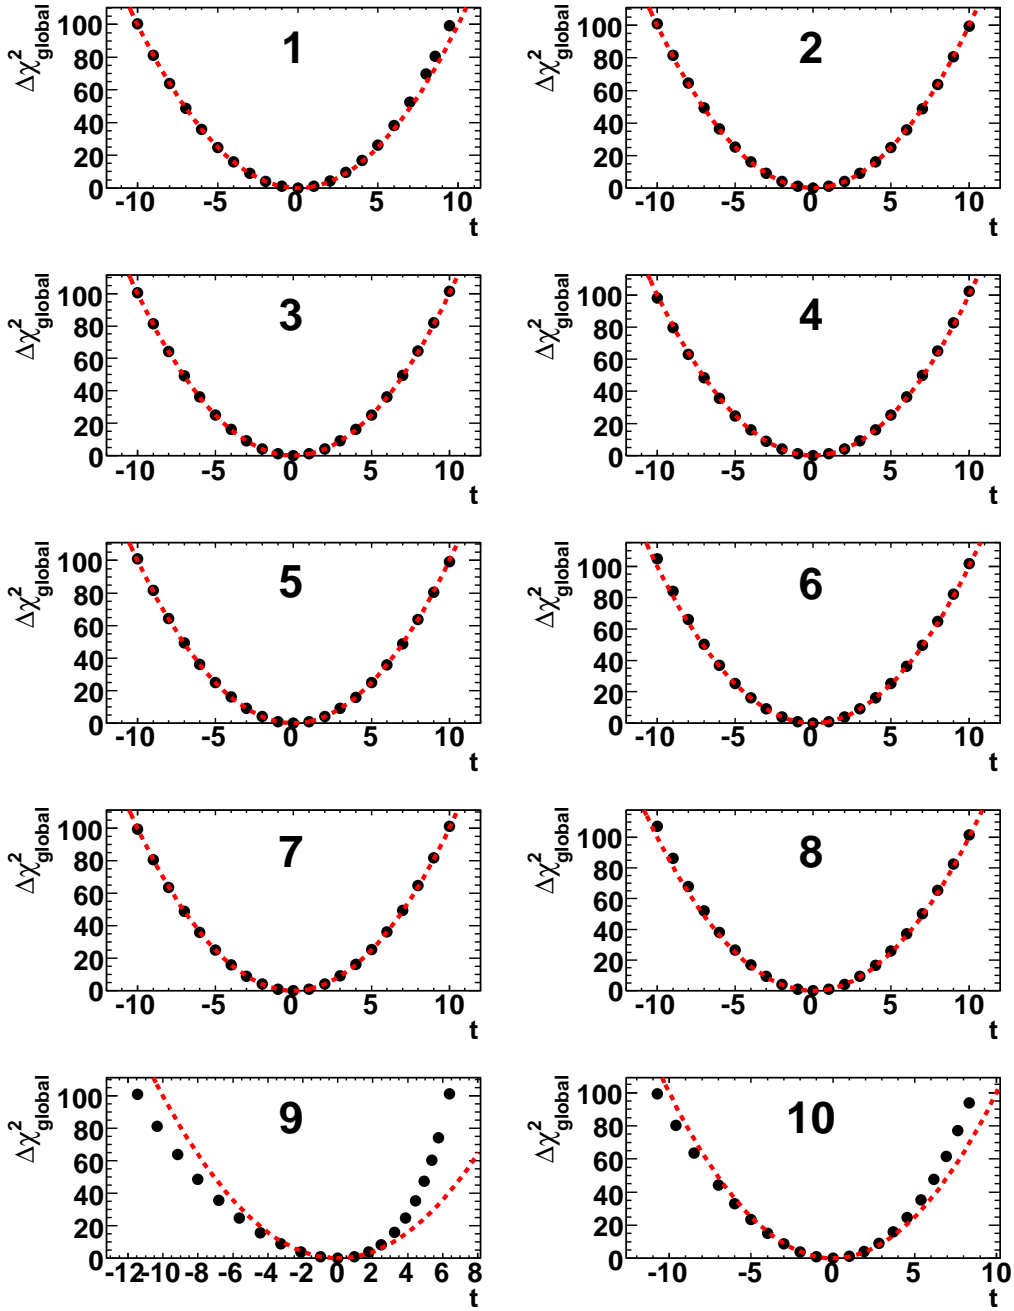


Figure 5: The points (\bullet) show $\Delta\chi^2_{\text{global}}$ as a function of the distance along each eigenvector direction, t , defined in (49), for eigenvectors numbered 1–10 corresponding to the 10 smallest eigenvalues. The dashed curve is the ideal case, $\Delta\chi^2_{\text{global}} = t^2$.

MSTW 2008 NLO PDF fit

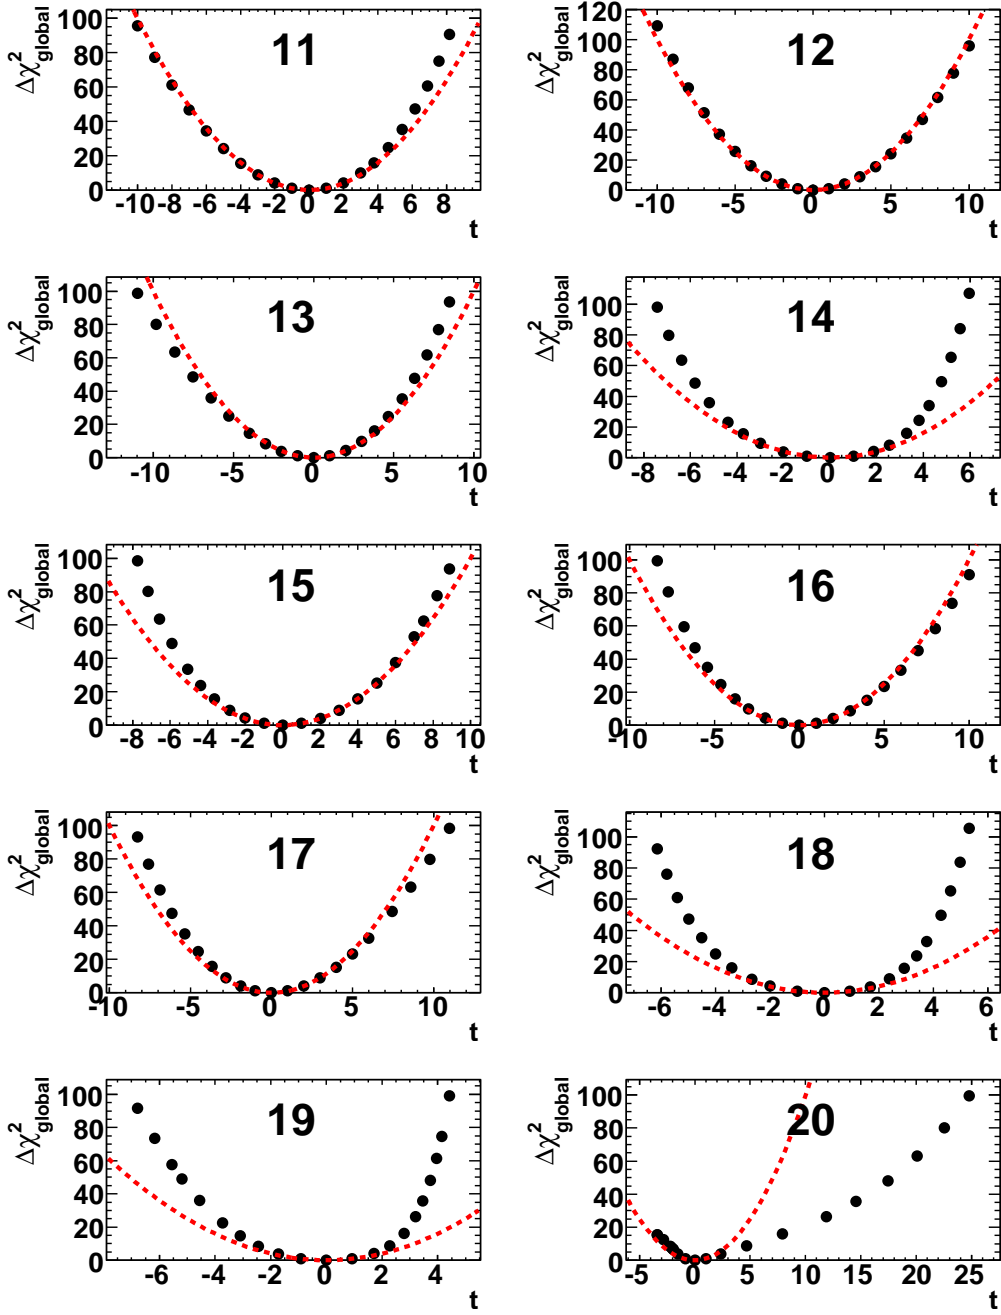


Figure 6: The points (\bullet) show $\Delta\chi^2_{\text{global}}$ as a function of the distance along each eigenvector direction, t , defined in (49), for eigenvectors numbered 11–20 corresponding to the 10 largest eigenvalues. The dashed curve is the ideal case, $\Delta\chi^2_{\text{global}} = t^2$.

certain amount of redundancy in parameters is a general feature of the full global fits obtained by CTEQ and MRST which have sufficient parametric freedom to account for all features of the data.

The behaviour of $\Delta\chi_{\text{global}}^2$ for the 20 eigenvectors in the current NLO fit is shown in Figs. 5 and 6. For the lowest 13 or so eigenvalues the quadratic approximation is very good. Above this we see some departure from quadratic behaviour, though there is still good agreement with the ideal case for relatively low t . The peculiar behaviour for eigenvector 20 is related to the fact that this is mainly concerned with the high- x part of $s_v = s - \bar{s}$ at input, and this will be discussed in Section 6.5.

6.2 Dynamic determination of tolerance

Ideally, with the standard “parameter-fitting” criterion, we would expect the errors to be given by the choice of tolerance $T = (\Delta\chi_{\text{global}}^2)^{1/2} = 1$ for the 68% (one-sigma) confidence level (C.L.) limit¹², or $T^2 = 2.71$ for the 90% C.L. limit [131, 132]. This is appropriate if fitting consistent data sets with ideal Gaussian errors to a well-defined theory. However, in practice, there are some inconsistencies between the independent fitted data sets, and unknown experimental and theoretical uncertainties, so the parameter-fitting criterion is not appropriate for global PDF analyses. One possible approach is to attempt to only include completely consistent data sets. This was tried in Ref. [136], but it was found that the ZEUS and NMC data were inconsistent with H1, BCDMS and E665, so only the latter three data sets were used in their fit. Clearly, this does not seem like a practical approach to follow.

So how can we choose the value of the tolerance T so as to obtain more reliable errors? Instead, we can appeal to the much weaker “hypothesis-testing” criterion, where the eigenvector PDF sets are treated as alternative hypotheses. Very roughly, a fit is judged to be “good” if each data set n , consisting of N data points, has $\chi_n^2 \simeq N \pm \sqrt{2N}$ [138]. More precisely, ranges of χ_n^2 corresponding to a 90% C.L. limit, for example, can be calculated, then the value of the tolerance $T = (\Delta\chi_{\text{global}}^2)^{1/2}$ can be chosen to ensure that each data set is described within its 90% C.L. limit. Assuming that χ_n^2 follows the χ^2 -distribution with N degrees of freedom, which has a probability density function

$$P_N(\chi^2) = \frac{(\chi^2)^{N/2-1} e^{-\chi^2/2}}{2^{N/2} \Gamma(N/2)}, \quad (56)$$

then the 90th percentile, ξ_{90} , is obtained by solving

$$\int_0^{\xi_{90}} d\chi^2 P_N(\chi^2) = 0.90. \quad (57)$$

Similarly for the 68th percentile, ξ_{68} , and for the most probable value, $\xi_{50} \simeq N$. In Fig. 7 we show (56) for $N = 163$ corresponding to a typical data set included in the global fit (BCDMS $\mu p F_2$). We have $\xi_{50} = 162.3$, $\xi_{68} - \xi_{50} = 8.7$ and $\xi_{90} - \xi_{50} = 24.2$, cf. $\sqrt{2N} = 18.1$.

¹²To be precise, rather than using exactly 68% we use $\text{erf}(1/\sqrt{2}) = 0.682689$.




Figure 7: Probability density function of the χ^2 -distribution with 163 degrees of freedom.

We now describe the new procedure we adopt in determining the tolerance for each eigenvector direction: it follows closely, but significantly extends, similar studies carried out earlier by the CTEQ group [126, 133, 134]. Since the new procedure is quite complicated, we will break the explanation down into steps, supplemented by example plots from the current NLO analysis.

- (i) First we plot the change in χ_n^2 for each data set, as the difference from the value at the *global* minimum, $\chi_{n,0}^2$, when moving along a particular eigenvector direction in units of $(\Delta\chi_{\text{global}}^2)^{1/2}$. An example for a subset of the most constraining data sets for eigenvector number 13 is shown in Fig. 8. The points (●) are generated for fixed values of the distance, $(\Delta\chi_{\text{global}}^2)^{1/2}$, between 0 and 10 in each “+” or “−” eigenvector direction. Note that we adjust t in (49) until the desired $(\Delta\chi_{\text{global}}^2)^{1/2}$ is obtained: these two quantities are equal only in the strict quadratic approximation. These points (●) are then fitted to a quadratic¹³ function shown by the solid lines. Note from Fig. 8 that the minimum value of χ_n^2 for a particular data set does not in general occur at the global minimum, indicating some tension between data sets included in the global fit, for example, between the E866/NuSea pp Drell–Yan data and the NuTeV and CHORUS νN xF_3 data.
- (ii) We define the 90% C.L. region for each data set by the condition that

$$\chi_n^2 < \left(\frac{\chi_{n,0}^2}{\xi_{50}} \right) \xi_{90}, \quad (58)$$

and similarly for the 68% C.L. region. Note that we define the 90% C.L. region (58) for χ_n^2 by renormalising ξ_{90} by a factor $\chi_{n,0}^2/\xi_{50}$ [133]. This procedure is necessary to take

¹³The one exception is for the very asymmetric χ_n^2 profile for the NuTeV dimuon data plotted against the distance along eigenvector number 20, associated with the η^- parameter controlling the high- x behaviour of $s_v = s - \bar{s}$, which we instead fit to a quartic function.

MSTW 2008 NLO PDF fit

Eigenvector number 13

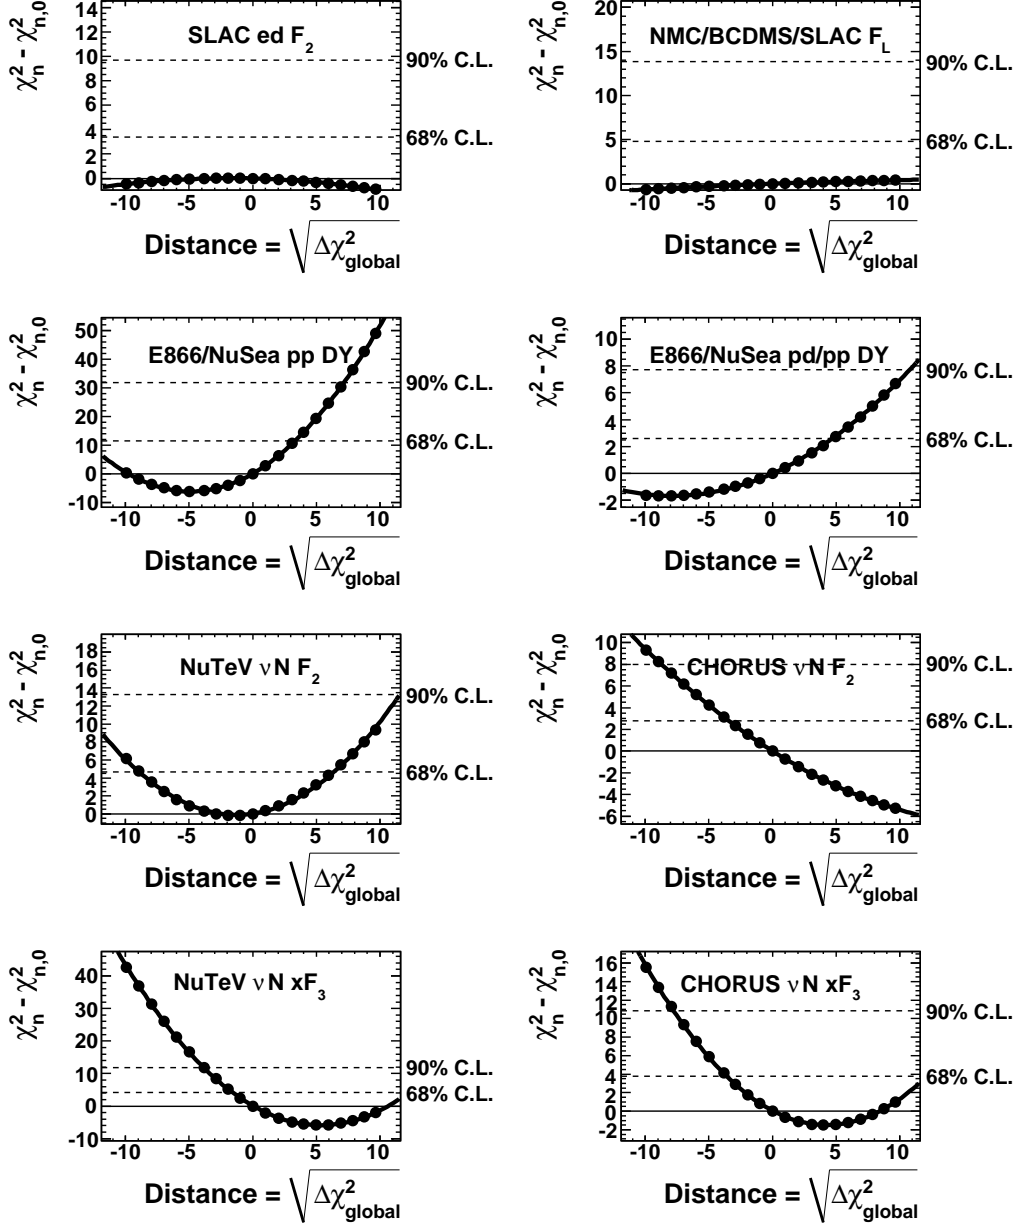


Figure 8: Change in χ_n^2 for a subset of data sets n , compared to the values at the global minimum ($\chi_{n,0}^2$), when moving along eigenvector number 13 in units of $(\Delta\chi_{\text{global}}^2)^{1/2}$. The 90% and 68% C.L. regions determined according to (58) are indicated.

account of the fact that the value of $\chi_{n,0}^2$ at the global minimum may be quite far from the most probable value of ξ_{50} of the particular data set n . For the example of the BCDMS $\mu p F_2$ data, the value of the χ_n^2 in our NLO fit is 182.2, which would lie outside the 68% C.L. region shown in Fig. 7 if the absolute values of χ_n^2 were used instead. After applying the rescaling factor of $\chi_{n,0}^2/\xi_{50} = 1.12$, the 90% C.L. region is given by $\chi_n^2 - \chi_{n,0}^2 < 27.2$ and the 68% C.L. region by $\chi_n^2 - \chi_{n,0}^2 < 9.8$. The 90% and 68% C.L. regions determined in this way are shown as the horizontal dashed lines in Fig. 8.

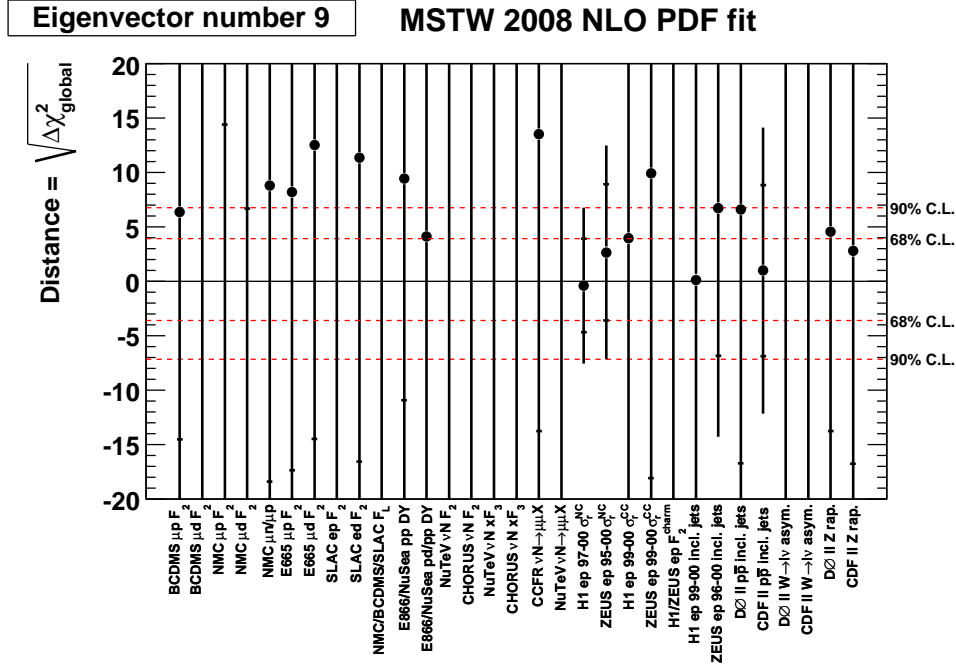
- (iii) For each eigenvector, we then record the value of the distance, $(\Delta\chi_{\text{global}}^2)^{1/2}$, for which the χ_n^2 for each data set is minimised, together with the 90% and 68% C.L. limits defined by the intercepts of the quadratic curves with the horizontal dashed lines in Fig. 8. These values are shown for eigenvectors 9 and 13 in Fig. 9, where the points (\bullet) indicate the values of the distance, $(\Delta\chi_{\text{global}}^2)^{1/2}$, for which χ_n^2 is minimised, while the inner error bars extend across the 68% C.L. region and the outer error bars extend across the 90% C.L. region defined by (58).
- (iv) For each of the eigenvectors, we choose the values of the tolerance $T = (\Delta\chi_{\text{global}}^2)^{1/2}$, indicated by the horizontal dashed lines in Fig. 9, so that all data sets are described within their 90% or 68% C.L. regions. For eigenvector 9 we see that the tolerance in the “+” direction is fixed by the H1 NC data and in the “−” direction by the ZEUS NC data, while for eigenvector 13 we see that the tolerance in the “+” direction is fixed by the E866/NuSea pp Drell–Yan data and in the “−” direction by the NuTeV $\nu N xF_3$ data.

We summarise the values of the tolerance obtained by this procedure in Fig. 10 for each of the 20 eigenvectors of the NLO fit. The inner error bars indicate the tolerance for a 68% C.L. limit while the outer error bars indicate the tolerance for a 90% C.L. limit. The labels placed at the end of the error bars indicate the name of the data set which fixes the value of the tolerance for a 90% C.L. limit. Note that the results shown for eigenvectors 9 and 13 correspond to the values that we extracted from the plots of Fig. 9.

The procedure just outlined is an extension of that previously employed by CTEQ and MRST. These previous analyses used a fixed value of the tolerance $T = (\Delta\chi_{\text{global}}^2)^{1/2}$ in producing eigenvector PDF sets, namely $T = \sqrt{100}$ for CTEQ and $T = \sqrt{50}$ for MRST. Analogous plots to those shown in Fig. 9 were originally presented by CTEQ in Appendix B.4 of Ref. [126]. Inspection of Fig. 10 shows that our new dynamic tolerance required to ensure that all data sets are described within their 90% C.L. limits are almost all in the region $T \sim \sqrt{50}$, which is close to the MRST value and suggests that the CTEQ tolerance ($T = \sqrt{100}$) is too large.¹⁴ However, an even smaller tolerance is obtained for some eigenvectors, in particular those associated with strange quarks and the $\bar{d} - \bar{u}$ difference, both of which are constrained by a small number of data sets generally each with a relatively small number of data points.

¹⁴Note, however, that recent CTEQ analyses [84, 87] apply unspecified weight factors to the data from certain experiments in their χ_{global}^2 definition, which in practice achieves a similar effect to our dynamic tolerance method.

(a)



(b)




Figure 9: Ranges of $(\Delta\chi^2_{\text{global}})^{1/2}$ along (a) eigenvector 9 and (b) eigenvector 13 for which data sets are satisfied within their 90% C.L. limit (outer error bars) or 68% C.L. limit (inner error bars). The points (•) indicate the minimum with respect to each particular data set. The tolerance, indicated by the horizontal dashed lines, is chosen to ensure that all data sets are described within their 68% or 90% C.L. limits defined by (58).

MSTW 2008 NLO PDF fit




Figure 10: Tolerance for each eigenvector direction determined dynamically from the criteria that each data set must be described within its 90% C.L. (Eq. (58)) (outer error bars) or 68% C.L. limit (inner error bars). The labels give the name of the data set which sets the 90% C.L. tolerance for each eigenvector direction.

6.3 Uncertainties on input PDFs

We use the values of the *dynamic* tolerance shown in Fig. 10 to generate the PDF eigenvector sets according to (49), which can be written as

$$a_i(S_k^\pm) = a_i^0 \pm t_k^\pm e_{ik}, \quad (59)$$

where t_k^\pm is adjusted to give T_k^\pm , with T_k^\pm the values shown in Fig. 10. We provide two different sets for each fit corresponding to either a 90% or 68% C.L. limit. Note that the ratio of the PDF uncertainties calculated using these two sets is not simply an overall factor of $\sqrt{2.71} = 1.64$, as it would be if choosing the tolerance according to the usual parameter-fitting criterion. Even in the simplest case, where the data set fixing the tolerance is the same for the 90% and 68% C.L. limits, and assuming linear error propagation, then the ratio of the T_k^\pm values would be $(\xi_{90} - \xi_{50})/(\xi_{68} - \xi_{50})$, which is a function of the number of data points N in the data set which fixes the tolerance, and takes a value around 1.7 for typical $N \sim 10$ –1000.

MSTW 2008 NLO PDF fit (68% C.L.)

Contribution to uncertainty from eigenvector number 9

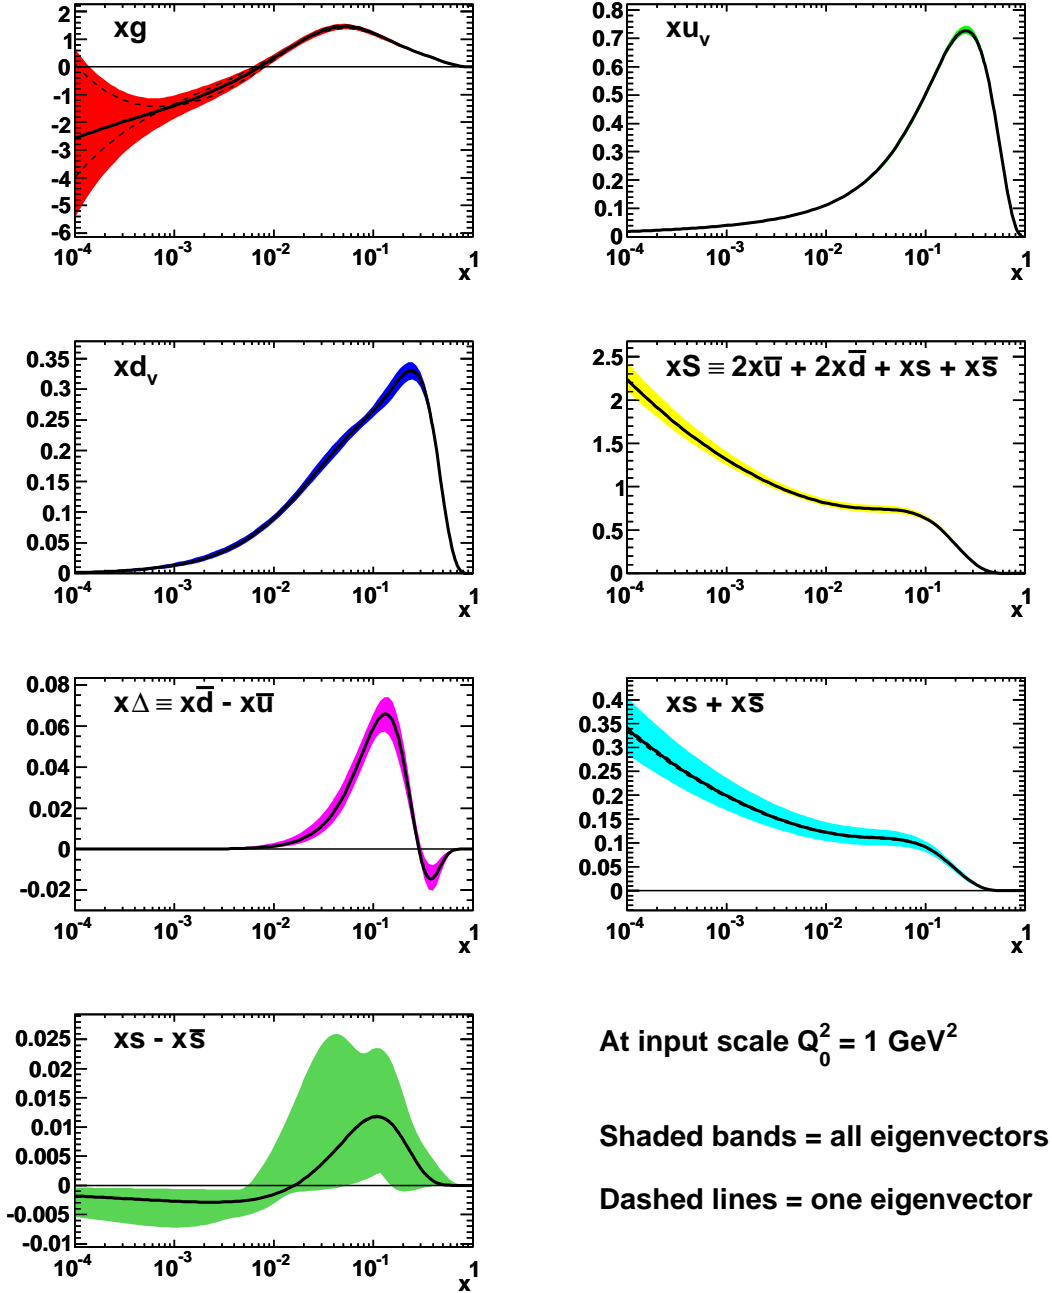


Figure 11: The total uncertainty on the input PDFs (shaded bands), given by (51) and (52), and the contribution from $k = 9$ in these equations (dashed lines). For this eigenvector, the dashed lines are hidden under the central solid lines with the exception of xg and $xs + x\bar{s}$.

The total one-sigma uncertainty on the input parton distributions at $Q_0^2 = 1 \text{ GeV}^2$, calculated using (51) and (52), is shown as the shaded bands in Fig. 11. Note that the somewhat peculiar shape of the uncertainty on $s - \bar{s}$ is clearly due to the limited freedom in the input parameterisation: only two free parameters controlling the normalisation and the power of $(1 - x)$.

Recall that each eigenvector \mathbf{v}_k is a linear combination of the input PDF parameter displacements $a_i - a_i^0$ and that the corresponding eigenvector PDF sets S_k^\pm are responsible for a certain uncertainty on the parton distributions. Sometimes an eigenvector corresponds to one particular input distribution, but often a single eigenvector contributes to the uncertainty on a variety of the input distributions. The dashed lines in Fig. 11 show the contribution to the uncertainty from a single eigenvector (in this case, number 9), i.e. only the $k = 9$ term in (51) and (52). Clearly, eigenvector 9 mainly contributes to the uncertainty on the input gluon distribution. This can be made clearer by plotting the fractional contribution to the uncertainty on the input PDFs from a single eigenvector, i.e. the contributions to (51) and (52) from a single value of k divided by the total uncertainties. The fractional contribution to the uncertainty from eigenvector number 9 is shown in Fig. 12 and that from eigenvector number 13 in Fig. 13.

A general summary of all the plots of the type shown in Figs. 12 and 13 is presented in Table 5. The bold entries show which eigenvectors k are mainly responsible for the uncertainty on the various input distributions. The three numbers in each entry correspond to (ten times) the fractional contribution to the uncertainties ($\Delta F_k / \Delta F$ for the $F = g, u_v, \dots$ input distributions, where ΔF_k is the contribution to (51) and (52) from a single eigenvector k) in the small x ($x < 0.01$), medium x ($0.01 < x < 0.1$) and large x ($x > 0.1$) regions, respectively, for the NLO global fit with 68% C.L. uncertainties. We see that each input distribution is sensitive to only a small number of eigenvectors. An indication of which data sets are particularly important for which flavour of parton distribution can be obtained by correlating the results of Fig. 10 with the table. For example, from Table 5 we see that the gluon distribution is particularly sensitive to eigenvectors 9, 11 and 19. The tolerance for each of these eigenvectors is fixed by the H1 and ZEUS NC data, apart from for the 68% C.L. tolerance for eigenvector 19 in the “ $-$ ” direction, where it is fixed by the CDF Run II high- p_T jet data.

In fact, the entries in Table 5 can be used in conjunction with the eigenvector plots, such as those used as examples in Fig. 9, to pinpoint which particular data sets constrain the various parton distributions. For example, the entries in Table 5 show that eigenvector 13 gives the biggest contribution to the high- x sea distribution uncertainty and is also very significant for the $\bar{d} - \bar{u}$ distribution at high x . Inspection of the plot of eigenvector 13 versus the data, Fig. 9(b), shows that the tightest constraints on the eigenvector 13 come from the data sets collected in the following experiments: E866/NuSea pp and pd/pp Drell–Yan, NuTeV and CHORUS $\nu N F_2$ and NuTeV and CHORUS $\nu N xF_3$. The former are the obvious direct constraint on the sea quarks, while the others influence the decomposition into the well-determined valence quarks and sea quarks. Eigenvector 13 also contains information on the d_v uncertainty, which is determined by the neutrino DIS data which weight down quarks more highly than neutral-current DIS on proton targets. Similarly, eigenvector 9 only influences the gluon uncertainty to any great degree, so it is no surprise that it is constrained by the H1 and ZEUS NC data (and CDF Run II high- p_T jet data); see Fig. 9(a). Surprisingly, the summary plot in Fig. 10 shows that the Tevatron jet data are never the main constraint on any eigenvector, apart from for the 68% C.L. tolerance for

MSTW 2008 NLO PDF fit (68% C.L.)

Fractional contribution to uncertainty from eigenvector number 9

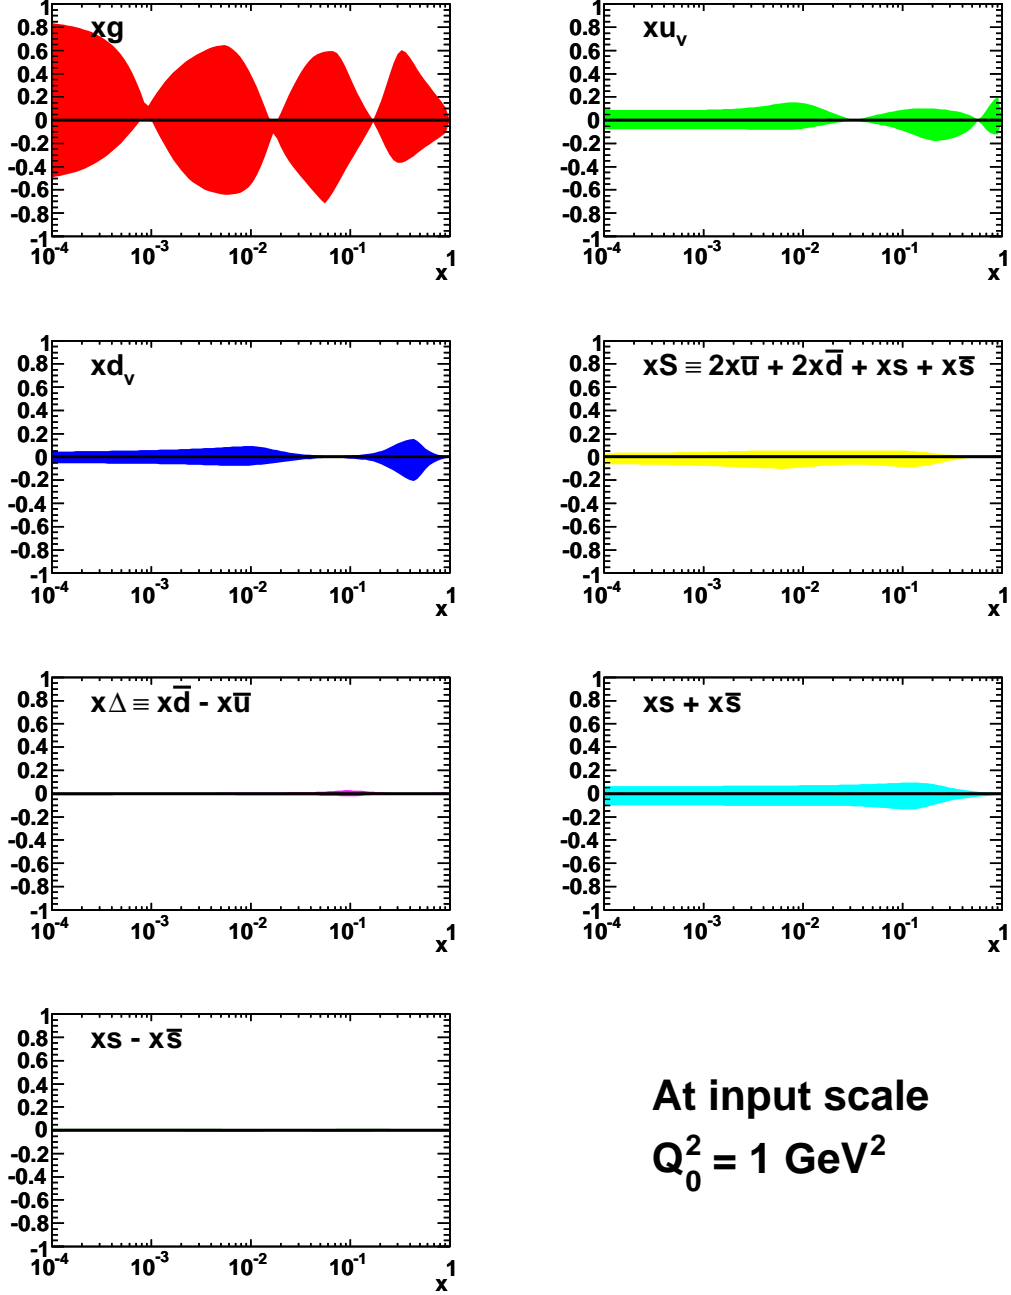


Figure 12: The fractional contribution to the uncertainty on the input PDFs from eigenvector 9, defined by the contributions to (51) and (52) from $k = 9$, divided by the total uncertainties.

MSTW 2008 NLO PDF fit (68% C.L.)

Fractional contribution to uncertainty from eigenvector number 13

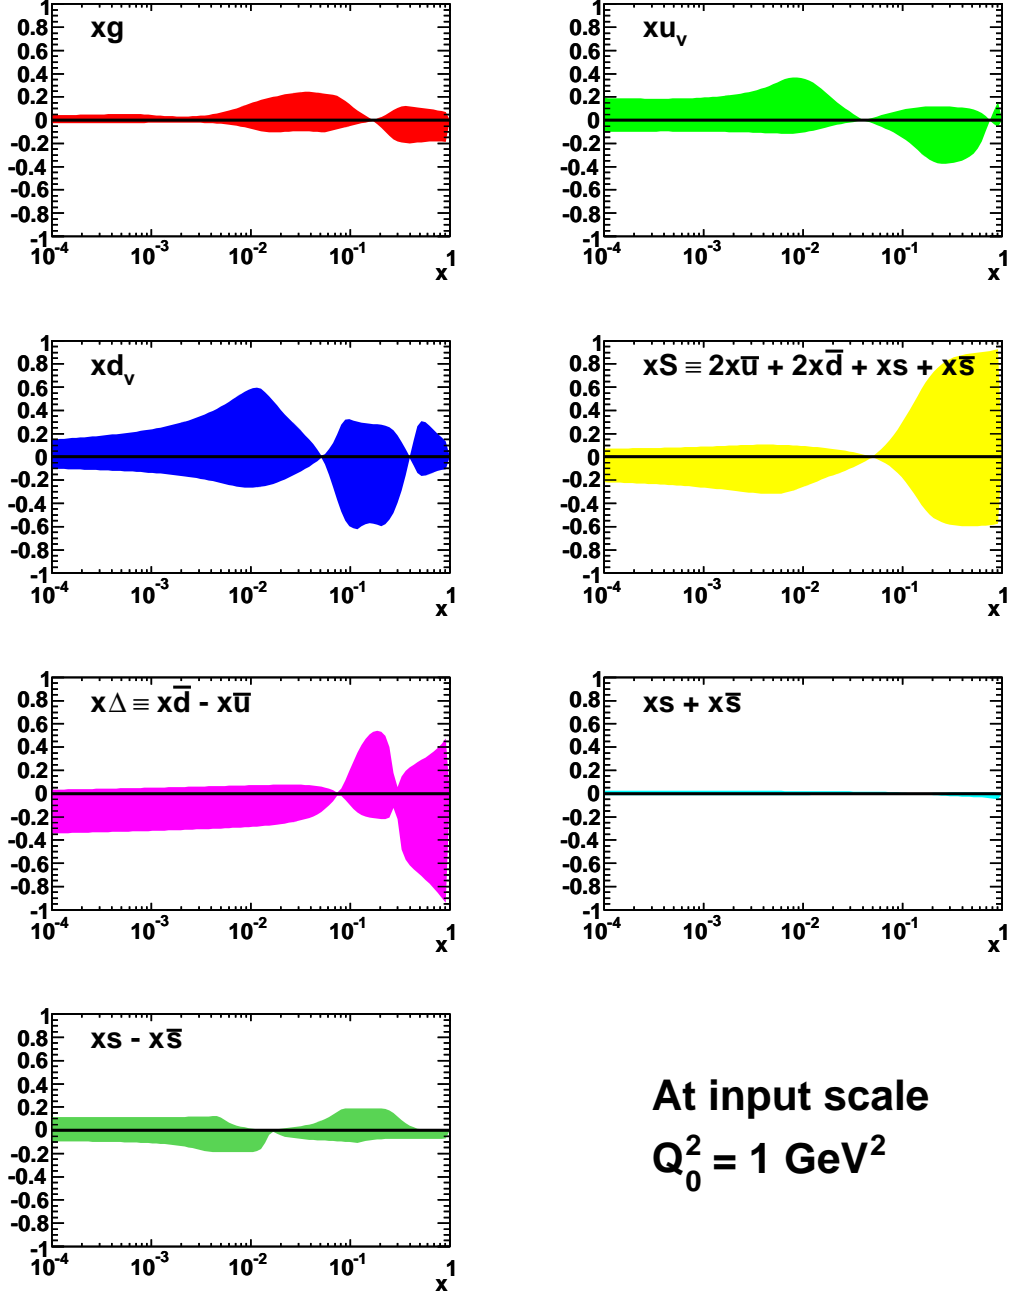


Figure 13: The fractional contribution to the uncertainty on the input PDFs from to eigenvector 13, defined by the contributions to (51) and (52) from $k = 13$, divided by the total uncertainties.

Eigenvector	g	u_v	d_v	$\bar{S}(\text{ea})$	$\bar{d} - \bar{u}$	$s + \bar{s}$	$s - \bar{s}$
1	5 2 1	—	—	—	—	—	—
2	1 0 0	3 3 3	—	1 1 1	—	2 3 3	—
3	0 1 0	1 2 2	—	—	—	4 4 5	—
4	—	0 1 1	—	0 1 1	1 5 6	2 2 3	—
5	1 1 1	1 1 2	—	2 2 2	0 0 1	1 1 2	4 4 4
6	3 4 3	3 3 5	—	7 7 7	0 2 2	4 6 6	9 9 9
7	1 1 1	—	4 4 6	1 1 1	0 0 1	—	1 0 0
8	1 2 1	4 4 6	1 2 2	—	0 2 2	—	—
9	8 7 6	1 1 2	1 1 2	1 1 1	—	1 1 1	—
10	2 1 1	0 1 3	2 5 6	2 2 1	1 2 2	—	1 1 1
11	8 6 8	7 6 7	8 8 8	6 6 5	4 4 2	4 5 5	—
12	2 4 4	1 1 4	2 2 2	1 1 2	10 9 8	1 1 1	—
13	1 2 2	4 4 4	6 6 6	3 2 9	3 3 9	—	2 2 2
14	2 2 1	1 1 4	9 9 9	1 1 1	2 4 3	0 1 0	—
15	7 4 3	5 4 5	3 4 2	7 7 5	3 3 2	2 4 2	—
16	0 1 1	2 1 1	2 2 2	2 2 3	0 1 7	4 4 6	1 1 1
17	1 1 1	1 1 0	1 2 2	2 0 2	1 1 5	7 6 8	1 1 1
18	1 2 2	8 9 4	2 2 5	1 2 1	0 5 5	1 0 1	1 1 1
19	6 8 9	4 5 4	1 1 1	2 3 3	1 2 2	2 2 3	0 0 1
20	1 1 0	2 4 2	2 4 4	1 2 6	0 2 4	1 2 4	9 10 10

Table 5: The three numbers in each entry are the fractional contributions to the uncertainties ($\Delta F_k/\Delta F$ for the $F = g, u_v, \dots$ input distributions, where ΔF_k is the contribution to (51) and (52) from a single eigenvector k) in the small x ($x < 0.01$), medium x ($0.01 < x < 0.1$) and large x ($x > 0.1$) regions, respectively, arising from eigenvector k in the NLO global fit, using 68% C.L. errors. Each number has been multiplied by ten; for example, 4 denotes 0.4. Major contributions are shown in bold type. For a precise value of x , and symmetric errors, the sum of the squares of each column should be 100. However, the entries shown are the maximum fraction in each interval of x , so often do not satisfy this condition.

eigenvector 19 in the “—” direction (not shown here). This does not necessarily mean that the Tevatron jet data provide *no* constraint, however, and this will be discussed in detail in Section 12.

A complete set of plots related to the error analysis and PDF uncertainties discussed in Sections 6.2 and 6.3, similar to the examples shown in Figs. 5–13, is available at Ref. [27]. These supplementary plots should prove valuable for the end users of the PDFs, in investigating the sensitivity of different observables to different eigenvector directions, and in finding the corresponding data sets which constrain those particular eigenvector directions.

6.4 Effect of free data set normalisations on PDF uncertainties

As already noted, one of our significant improvements in the uncertainty determination in the present analysis is the inclusion of the data normalisation errors while determining the PDF uncertainties. In previous MRST (and CTEQ) fits these have been held fixed during the error propagation. To account for free data set normalisations, we first calculate the 39×39 covariance matrix with an additional 19 free parameters¹⁵ corresponding to data set normalisations, in addition to the usual 20 free input PDF parameters. In principle the inclusion of free data set normalisations in PDF uncertainties can be done in two different ways.

The simplest way is to diagonalise the 39×39 covariance matrix to obtain 39 eigenvectors and to determine the tolerance values dynamically as described above. However, this method involves a large number of eigenvectors, many of which actually consist largely of the variation of data set normalisations. It is clearly more physically intuitive to use eigenvectors corresponding only to input PDF parameters.

An alternative method, which has this advantage and which we use in the present analysis, is to remove the rows and columns from the 39×39 covariance matrix corresponding to data set normalisations, leaving only a 20×20 covariance matrix corresponding to input PDF parameters, which contains, however, information on free data set normalisations from the inversion of the Hessian matrix [131]. Then the 20×20 covariance matrix can be diagonalised to give 20 eigenvectors, and the tolerance values and eigenvector PDF sets produced as described above, with the modification that every time the PDF parameters are changed, when moving along an eigenvector direction, the χ_{global}^2 should be minimised with respect to the data set normalisations.¹⁶

In the limit of quadratic behaviour for the χ_{global}^2 function and equal tolerance for each eigenvector we would obtain the same results for PDF uncertainties in each case. In Fig. 14 we show the effect of allowing data set normalisations to go free according to these two methods, compared to the case where data set normalisations are fixed as in the MRST analyses. The new procedure leads to an increase in uncertainties, but in most places this increase is very small. The exception is those parton distributions where the uncertainty is smallest, e.g. the up valence quark for $x \gtrsim 0.05$ and the light sea for x between 0.001 and 0.01. The absolute increase in uncertainty is 1% or less, but in these cases it is a significant contribution to the total. We see that the approach using data set normalisations in eigenvectors gives very similar results, demonstrating that despite our quartic penalty term for normalisations there is little breakdown of the quadratic approximation. Indeed, the uncertainty on normalisations determined from the fit and applying the tolerance approach is usually rather similar to the quoted experimental uncertainty. Hence, we do not actually encounter large penalties from varying normalisations, as we rarely vary by much more than the experimental one-sigma uncertainty.

¹⁵To be precise, there are 16 free parameters associated with data set normalisations together with the 3 parameters (r_1, r_2, r_3) associated with nuclear corrections (73). We do not allow the normalisation of the neutrino data to go free because it is very highly correlated with r_1 .

¹⁶Similarly, the χ_{global}^2 should be minimised with respect to the shifts in correlated systematic errors, r_k in Eq. (38), each time the PDF parameters are changed, but this is done automatically by using Eq. (39).

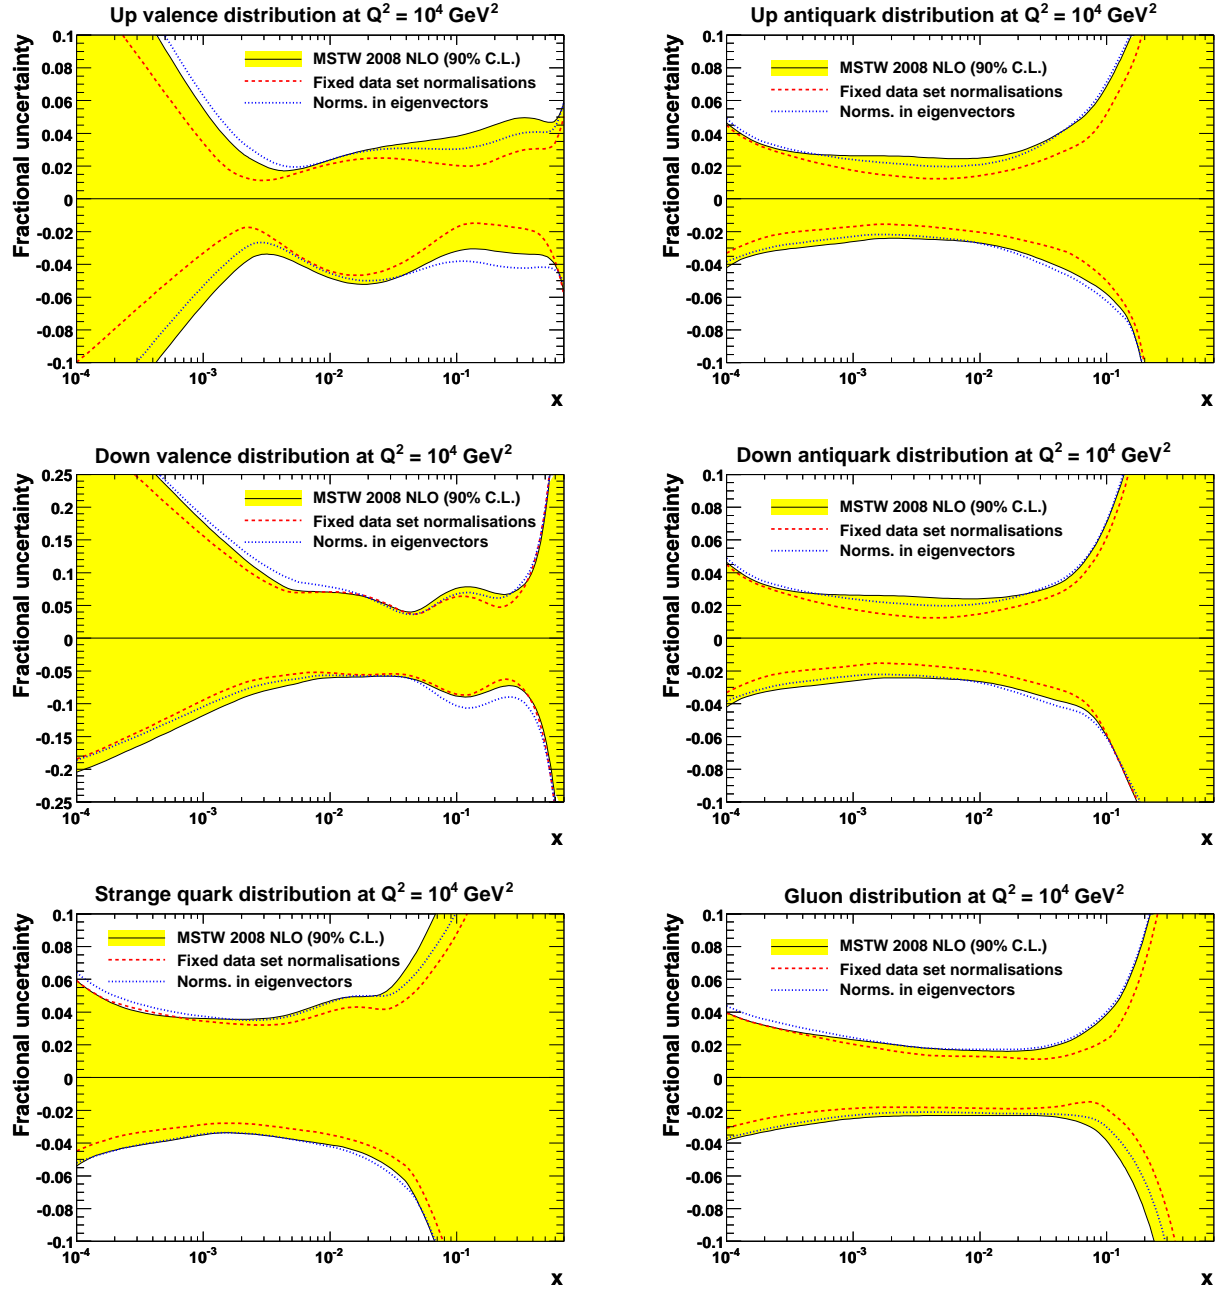


Figure 14: The fractional uncertainty of the MSTW 2008 NLO PDFs at $Q^2 = 10^4 \text{ GeV}^2$ using 20 eigenvectors and allowing for free data set normalisations. We also show the fractional uncertainty obtained using 20 eigenvectors with fixed data set normalisations (as in the MRST analyses) and with 39 eigenvectors including 19 parameters associated with data set normalisations in addition to the 20 input PDF parameters.

6.5 Input parameterisation and uncertainties

There is a strong relationship between the input PDF parameterisation and the uncertainties which will be obtained. Too restrictive a parameterisation can limit the calculated uncertainty in an artificial manner. The parameterisation for the input PDFs in our analysis were presented in Section 3, specifically in Eqs. (6)–(12). The 28 free PDF parameters listed there allow a large degree of flexibility, but are sufficiently few that only a few turning points are possible for each distribution, presumably the case in reality for quantities closely related to physical variables. Indeed, as outlined in Section 6.1, many of the 28 free PDF parameters turn out to be very strongly (anti)correlated, and only 20 parameters remain after removing all combinations where this is the case, whose uncertainties are quoted in Table 4. This is 5 more parameters than in the MRST error sets [16, 21]. Here we briefly discuss the truly free parameters for each parton flavour, i.e. those parameters allowed to go free when calculating the covariance matrix used for error propagation, and outline the resulting effect on the uncertainty.

Valence quarks

For the up and down valence quark distributions we use the standard “MRS-type” parameterisation,

$$xv(x, Q_0^2) = A_v x^{\delta_v} (1-x)^{\eta_v} (1 + \epsilon_v x^{0.5} + \gamma_v x), \quad (60)$$

where the normalisation parameter A_v is determined from the number sum rules (14). In both cases the 3 parameters left free in the determination of the eigenvectors are δ_v , η_v and ϵ_v . These correspond to three reasonably independent regions of x . At high x the uncertainty is determined very largely by η_v . In detail,

$$\begin{aligned} v \pm \Delta v &\sim A_v (1 + \epsilon_v + \gamma_v) (1-x)^{\eta_v \mp \Delta\eta_v} \\ &\sim v (1-x)^{\mp \Delta\eta_v} \\ &\sim v [1 \mp \Delta\eta_v \ln(1-x)], \end{aligned} \quad (61)$$

where the last line applies if $|\Delta\eta_v \ln(1-x)|$ is small enough for the linear expansion of the exponential to be valid. Hence, the parameterisation allows the uncertainty to become very large as $x \rightarrow 1$, as required by the lack of data in this region. It is not as flexible as some of the competing parameterisations, e.g. the CTEQ6M PDFs [126]. However, it does limit the possibility of unusual shapes at very high x , or arguably undesirable features, e.g. the gluon becoming the hardest distribution at very high x ($x > 0.9$), as found in the CTEQ6.1M set. At small x the uncertainty is determined very largely by δ_v . Again in more detail,

$$\begin{aligned} x(v \pm \Delta v) &\sim A_v x^{\delta_v \mp \Delta\delta_v} \\ &\sim xv x^{\mp \Delta\delta_v} \\ &\sim xv [1 \pm \Delta\delta_v \ln(1/x)], \end{aligned} \quad (62)$$

where again the last line applies if $|\Delta\delta_v \ln(1/x)|$ is small enough for the linear expansion of the exponential to be valid. Similarly to the case at high x , the uncertainty can become very large

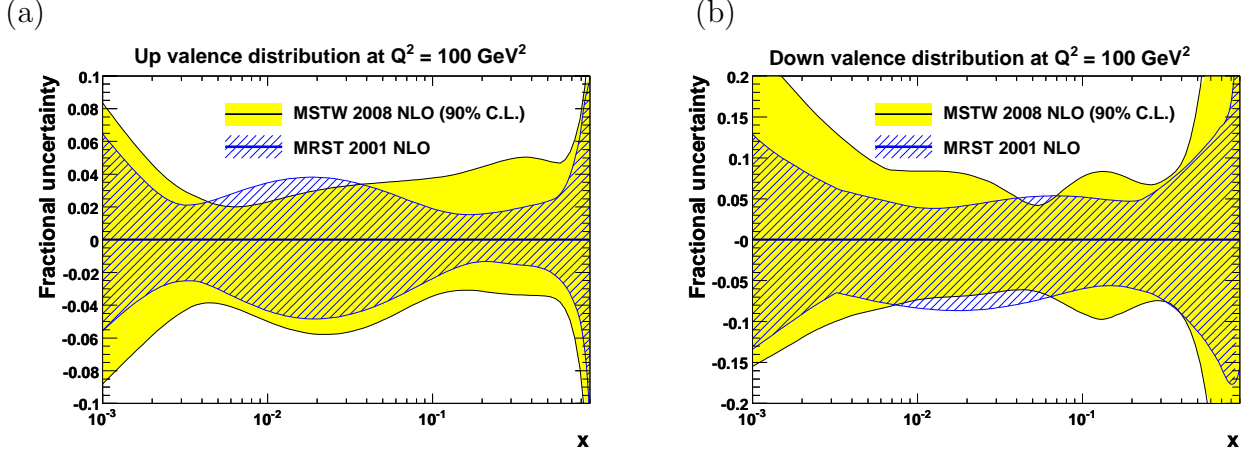


Figure 15: The fractional uncertainty on the (a) up valence quark and (b) down valence quark distributions at $Q^2 = 100 \text{ GeV}^2$ from the MSTW 2008 NLO fit compared to that from the MRST 2001 NLO fit [16].

at small x , but once the behaviour has been determined by the data at fairly low x , fixing the one parameter, even lower x values are simply an extrapolation. The third free parameter, ϵ_v , controls the variation of the input distribution at intermediate values of x (γ_v is more strongly correlated to η_v). In practice the uncertainties at very high x , and even more so at very small x , are also constrained by the sum rules. In the region where valence PDFs are constrained by data, i.e. $x = 0.01\text{--}0.75$, about 75% of the total number of quarks are found, so the number found in the extrapolation regions is relatively small and very well determined.

The fractional uncertainty for the up valence distribution at NLO is shown in Fig. 15(a) along with that from the MRST 2001 analysis [16]. In this case the parameters left free when determining the eigenvectors are the same and the uncertainties are larger mainly because of the fact that data set normalisations were fixed in the MRST analysis; see Fig. 14. The fractional uncertainty on u_v increases rapidly for very high x and also for $x \lesssim 0.003$. In the latter case the data constraint for $x \lesssim 0.01$ at lower Q^2 has moved to a smaller x due to evolution. The uncertainty for the down valence distribution is shown in Fig. 15(b). The free parameters in the MRST 2001 analysis were η_d , ϵ_d and γ_d . Hence, the MRST 2001 input parameterisation allowed less flexibility in the d_v uncertainty at small x and from the sum rule this fed in to the unconstrained high- x region also. This is clearly illustrated in Fig. 15(b).

Gluon

The gluon distribution is the most complicated case for PDF uncertainties and parameterisations. The input gluon in the original MRST fits [12, 13] had the same form of parameterisation as Eq. (60), with A_g determined by the momentum sum rule. However, with the advent of improved HERA data included in the MRST 2001 fit [14], it was noticed that if the evolution is started from $Q_0^2 = 1 \text{ GeV}^2$ the input gluon preferred to be negative at very small x and an additional

term $A_{g'} x^{\delta_{g'}} (1-x)^{\eta_{g'}}$ was added to facilitate this. To this date all competing parameterisations, apart from the very recent NNPDF1.0 set [137]¹⁷, have a gluon which at small x behaves like $xg(x, Q_0^2) \sim x^{\delta_g}$, i.e. is controlled by a single power. This means that

$$g \pm \Delta g \sim g [1 \pm \Delta \delta_g \ln(1/x)], \quad (63)$$

i.e. the uncertainty grows linearly with $\ln(1/x)$ and there is no scope for a rapidly expanding uncertainty as data constraints run out. This is much more of an issue for the gluon than for valence quarks, as the momentum sum rule offers a far less direct constraint than the number sum rules as $x \rightarrow 0$. However, there is another complication to consider, namely

$$\Delta g(x, Q_0^2) \sim g(x, Q_0^2) \Delta \delta_g \ln(1/x), \quad (64)$$

and so as $g(x, Q_0^2)$ becomes smaller then so does $\Delta g(x, Q_0^2)$. If $g(x, Q_0^2)$ is very small, then the absolute input uncertainty for the gluon is very small, and at higher Q^2 the uncertainty is therefore determined entirely by evolution from higher- x , i.e. by the region where the gluon distribution is better determined. Most PDF fitting groups find that $xg(x, Q^2)$ is indeed small at low Q^2 and small x . In this region the MRST (since 2001) and MSTW gluon distributions have the form,

$$xg(x, Q_0^2) = xg_1(x, Q_0^2) + xg_2(x, Q_0^2) \sim A_g x^{\delta_g} + A_{g'} x^{\delta_{g'}}, \quad (65)$$

which is more flexible than a single power. Not only does it allow the gluon to become negative at very small x , but it is also particularly important for the uncertainty,

$$\Delta g(x, Q_0^2) \sim \pm g_1(x, Q_0^2) \Delta \delta_g \ln(1/x) \pm g_2(x, Q_0^2) \Delta \delta_{g'} \ln(1/x), \quad (66)$$

where g_1 and g_2 represent the two independent terms in the gluon parameterisation. The interplay between the two terms allows for a large uncertainty at $x \lesssim 10^{-4}$ where the data constraint, from the Q^2 dependence of $F_2(x, Q^2)$ at HERA, diminishes rapidly.

A comparison of the fractional uncertainty for the present MSTW, CTEQ6.6 [87], Alekhin [139] and NNPDF1.0 [137] NLO gluon distributions is shown in Fig. 16. The fractional uncertainty for MSTW blows up very quickly at small x , whereas that for the Alekhin fit gets bigger relatively slowly (the same would be qualitatively true of the H1 [109] and ZEUS [140] fits), while that for CTEQ saturates, or can even decrease slightly, completely contrary to what is expected from the degree of constraint from data. These results are related to the input forms and values of Q_0^2 , and can be understood fairly easily. MRST/MSTW parameterise the starting distributions at $Q_0^2 = 1 \text{ GeV}^2$, and since we allow both positive and negative small- x contributions the uncertainty can be very large. We believe that this represents the true uncertainty at low x . The Alekhin, H1 and ZEUS gluon distributions are input at a higher scale. They behave like $x^{-\lambda}$ at small x , so the uncertainty is due to the uncertainty in one parameter, as shown in (63). Recent CTEQ gluon distributions are input at $Q_0^2 = 1.69 \text{ GeV}^2$, and behave like x^λ at small x where λ is large and positive, i.e. the input gluon is valence-like.¹⁸ This requires fine tuning: when evolving

¹⁷Despite the NNPDF1.0 input parameterisation formally behaving as $xg \sim x^{-0.2} \text{NN}_g(x)$, the neural network parameterisation $\text{NN}_g(x)$ should be sufficiently flexible that there is no dependence on the factor $x^{-0.2}$ [137].

¹⁸A study of the input gluon parameterisation by CTEQ was made in Appendix D of Ref. [126], but it did not address the impact of the input parameterisation on the uncertainties.




Figure 16: A comparison of the fractional uncertainty for the present MSTW, CTEQ6.6 [87], Alekhin [139] and NNPDF1.0 [137] NLO gluon distributions at $Q^2 = 5 \text{ GeV}^2$. All uncertainty bands represent a 90% C.L. limit.

backwards from a steep gluon at higher scales a valence-like gluon only exists for a very narrow range of Q^2 (if at all). In this case the small- x input gluon distribution is tiny, so there is a very small absolute error, as seen from (64). At higher Q^2 all the uncertainty is due to evolution and driven by the higher x and more well-determined gluon distribution. Hence the very small x gluon is no more uncertain than at $x = 0.001\text{--}0.01$.¹⁹ This does not seem to be realistic to us. The NNPDF1.0 [137] uncertainty shown in Fig. 16, which is obtained using a very different approach for the input parameterisation and error propagation, will be discussed in Section 14.

The above discussion of the input gluon parameterisation applies to all NLO and NNLO MRST/MSTW sets since 2001. However, in the MSTW 2008 analysis we have one more free parameter in the uncertainty determination, i.e. as well as δ_g , $\delta_{g'}$ and η_g we also have $\eta_{g'}$. The latter contributes to the uncertainty at intermediate x since its central value is much greater than that of η_g . The comparison of the uncertainty in the gluon distribution to that from the MRST 2001 and 2006 analyses is shown in Fig. 17. As well as increased uncertainty at small x one sees that the “neck” at $x \approx 0.13$ has disappeared in the new analysis.

Note that the fractional uncertainty for the gluon distribution from the MRST 2006 NNLO set, shown in Fig. 17(b), has an additional narrowing or “neck” at $x = 0.7$. This has a simple explanation. The MRST 2004 [18] and 2006 [21] PDF sets defined the input $\overline{\text{MS}}$ gluon parameterisation via a transformation from the DIS to the $\overline{\text{MS}}$ scheme, where the input gluon

¹⁹The same features exist for the “dynamical parton distributions” of Ref. [141]. However, in that case the starting scale is $Q_0^2 = 0.5 \text{ GeV}^2$, and the PDFs do not fit data as well as the “standard” PDFs which are obtained from evolution at a higher starting scale.

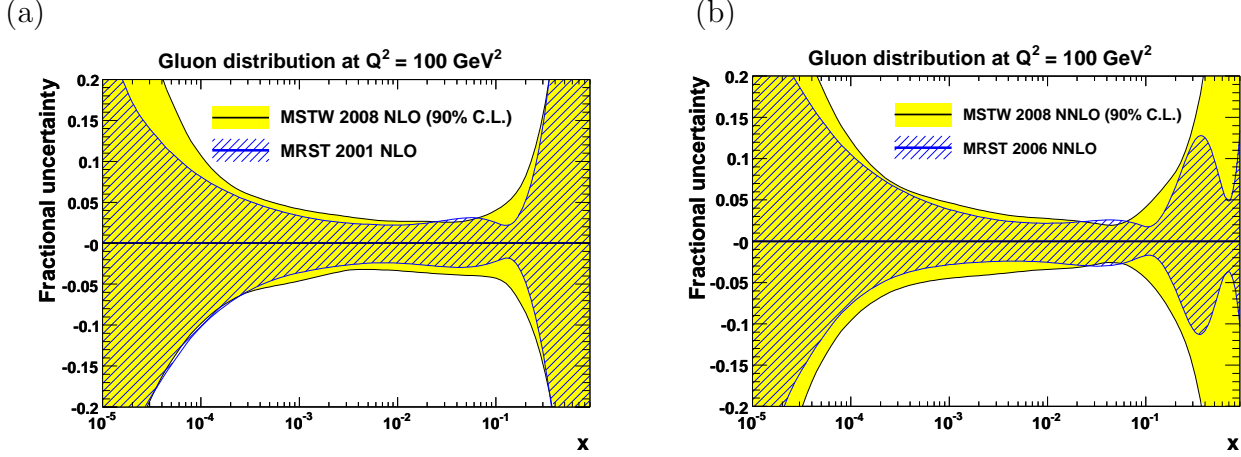


Figure 17: The fractional uncertainty on the gluon distribution at $Q^2 = 100 \text{ GeV}^2$ (a) from the MSTW 2008 NLO fit compared to that from the MRST 2001 NLO fit [16] and (b) from the MSTW 2008 NNLO fit compared to that from the MRST 2006 NNLO fit [21].

distribution in the DIS scheme was given by the standard parameterisation (10), giving an enhancement of the $\overline{\text{MS}}$ gluon at high x originating from the high- x quarks. Hence, in practice, at high x , i.e. above $x \approx 0.4$ – 0.5 , the MRST 2006 gluon distribution was dominated by the quark distributions via the $\text{DIS} \rightarrow \overline{\text{MS}}$ transformation. This meant that for $x \gtrsim 0.6$ the relative uncertainty on the gluon effectively became that on the quarks, which is much smaller, but does start diverging again at even higher x , as can clearly be seen in Fig. 17(b). As we will see in Section 12, the high- x enhancement in the gluon distribution is disfavoured by Tevatron Run II jet data, and so the present MSTW 2008 fits revert to the standard parameterisation (10) for the input gluon distribution directly in the $\overline{\text{MS}}$ scheme. Hence the uncertainty simply expands as $x \rightarrow 1$, as for the other PDFs.

The LO gluon distribution has no tendency to turn over at small x , so the second term in (10) is redundant. In this case we use the standard “MRS-type” parameterisation (60), but with all 4 parameters ($\delta_g, \eta_g, \epsilon_g, \gamma_g$) free in determining the eigenvectors.

Sea asymmetry, $\bar{d} - \bar{u}$

The sea asymmetry $\bar{d} - \bar{u}$ is parameterised at input as

$$x(\bar{d} - \bar{u})(x, Q_0^2) = A_{(\bar{d}-\bar{u})} x^{\eta_{(\bar{d}-\bar{u})}} (1-x)^{\eta_S+2} (1 + \gamma_{(\bar{d}-\bar{u})} x + \delta_{(\bar{d}-\bar{u})} x^2), \quad (67)$$

and the parameters contributing to the eigenvectors are the same as in the MRST 2001 fit. The distribution is constrained by data for $0.01 \lesssim x \lesssim 0.2$. At small x the fact that we choose a simple power means that once the data for $0.01 \lesssim x \lesssim 0.1$ cause the distribution to head to zero there is no scope for it to do otherwise at smaller x . We see no reason to suppose that, as a nonsinglet quantity, it should not go to zero as x goes to zero, due to considerations from Regge theory. At high x we choose the $(1-x)$ power to be $\eta_S + 2$, since $\bar{d}(x, Q^2) - \bar{u}(x, Q^2)$

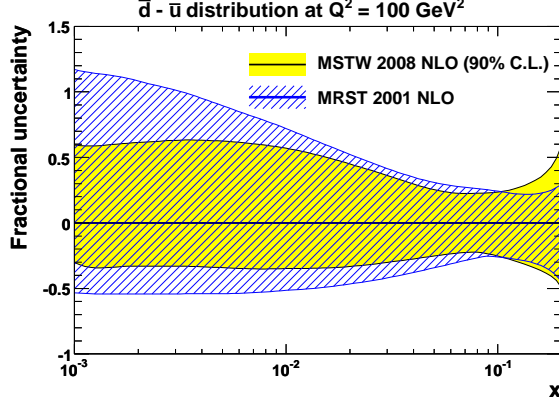


Figure 18: The fractional uncertainty on the $\bar{d} - \bar{u}$ distribution at $Q^2 = 100 \text{ GeV}^2$ from the MSTW 2008 NLO fit compared to that from the MRST 2001 NLO fit [16].

is becoming very small, and we attempt to constrain $\bar{u}, \bar{d} \geq 0$. Applying this constraint on the power of $(1-x)$ we then need the x^2 term instead of the more usual \sqrt{x} term to give sufficient flexibility at large x . The 3 eigenvector parameters control the normalisation, high x and low x . In fact, in some PDF sets one does then get negative central values of the $\bar{d}(x, Q^2)$ at the highest x , but at a level very small compared to the uncertainties. The uncertainties for $\bar{d} - \bar{u}$ in the MSTW 2008 PDF set are slightly smaller than for MRST 2001, as seen in Fig. 18, but this is due to the inclusion of new data and also smaller values of the tolerance.

Light sea quarks

The sea distribution $S \equiv 2\bar{u} + 2\bar{d} + s + \bar{s}$ is parameterised by the standard “MRS-type” form,

$$xS(x, Q_0^2) = A_S x^{\delta_S} (1-x)^{\eta_S} (1 + \epsilon_S \sqrt{x} + \gamma_S x), \quad (68)$$

In this case there is no sum rule constraint, so in principle there are 5 free parameters. Again, only 3 can contribute to the eigenvectors due to very large correlations. We would expect these to be δ_S , η_S and ϵ_S , but in practice δ_S has too large correlations with the gluon parameters, so we take A_S instead. Hence, in this case, at small x the uncertainty at input is due to A_S , and the uncertainty $\Delta S(x, Q_0^2)$ is simply proportional to $S(x, Q_0^2)$. However, at higher Q^2 the uncertainty is controlled by the gluon evolution which leads very quickly to a much larger uncertainty at very small x . This feature is shown in Fig. 19(a). Perhaps we underestimate the sea uncertainty at small Q^2 and very small x with this choice of free parameters. However, the uncertainty does follow the data constraint, i.e. $F_2(x, Q^2)$ constrains the quark distributions down to $x = 10^{-5}$ quite precisely but only for a narrow range of small Q^2 . In Fig. 19(b) we compare the sea distribution at $Q^2 = 100 \text{ GeV}^2$ to that from the MRST 2001 analysis. The uncertainties are larger in the new analysis mainly due to the extra freedom in the strange quark parameterisation, which feeds into the uncertainties for \bar{u} and \bar{d} , and also due to allowing data set normalisations to go free; see Fig. 14.

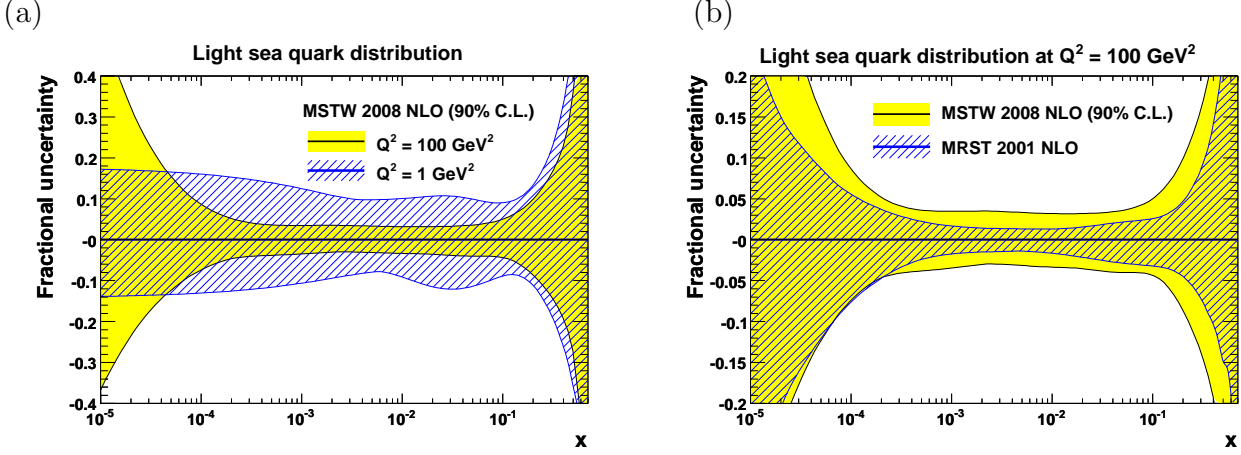


Figure 19: The fractional uncertainty on the light sea quark distribution $S \equiv 2\bar{u} + 2\bar{d} + s + \bar{s}$ at $Q^2 = 100 \text{ GeV}^2$ from the MSTW 2008 NLO fit compared to (a) the input sea quark distribution at $Q^2 = 1 \text{ GeV}^2$ and (b) the sea quark distribution from the MRST 2001 NLO fit [16].

Strange quarks

As mentioned before we now parameterise the strange quarks separately rather than assume that

$$s(x, Q_0^2) + \bar{s}(x, Q_0^2) = \kappa [\bar{u}(x, Q_0^2) + \bar{d}(x, Q_0^2)], \quad (69)$$

with $\kappa \approx 0.4\text{--}0.5$. We could use a completely free parameterisation, but the data only exist for $x \gtrsim 0.01$ so this would lead to an enormous uncertainty for $x \lesssim 0.01$. This is essentially what is done for the CTEQ6.6 distributions [87], but it is questionable whether this is realistic. As a rough constraint, for example, Regge trajectory considerations suggest that all flavour distributions have the same power as $x \rightarrow 0$, and we certainly think that it is sensible to impose this. However, we apply more detailed constraints from perturbative QCD reasoning. The strange quark has some non-insignificant mass, and we assume that it is this which qualitatively leads to the suppression compared to up and down sea quarks. This assumption gives us a handle on the type of parameterisation to use.

When charm and bottom quarks begin their evolution they evolve like massless quarks, but always lag behind. This leads to some suppression at all x for finite Q^2 . But at small x this suppression is roughly just a normalisation suppression, i.e. evolution makes the small- x powers δ_i the same. Indeed, other than the overall normalisation suppression, further effects are mainly at high x . This reasoning leads to us choosing only the normalisation A_+ and the high- x power η_+ as free parameters for $s(x, Q_0^2) + \bar{s}(x, Q_0^2)$. Indeed, effectively we choose a parameterisation of the form

$$\begin{aligned} xs(x, Q_0^2) + x\bar{s}(x, Q_0^2) &= A_+ x^{\delta_S} (1-x)^{\eta_+} (1 + \epsilon_S \sqrt{x} + \gamma_S x) \\ &\equiv \tilde{A} (1-x)^{\tilde{\eta}} xS(x, Q_0^2), \end{aligned} \quad (70)$$

where $\tilde{A} \equiv A_+/A_S$, $\tilde{\eta} \equiv \eta_+ - \eta_S$ and the small- x power is fixed to δ_S . Introducing distinct ϵ and

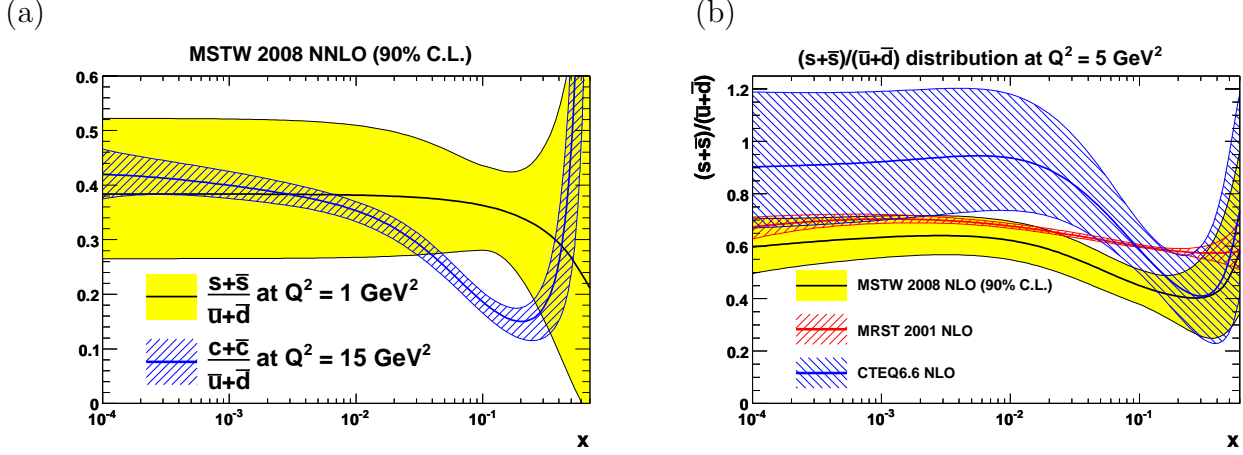


Figure 20: (a) $s + \bar{s}$ at $Q^2 = 1 \text{ GeV}^2$ compared to $c + \bar{c}$ at $Q^2 = 15 \text{ GeV}^2$, both as a fraction of $\bar{u} + \bar{d}$. (b) $(s + \bar{s})/(\bar{u} + \bar{d})$ at $Q^2 = 5 \text{ GeV}^2$ from the current NLO analysis compared to the same quantity from the MRST 2001 [16] and CTEQ6.6 [87] analyses.

γ parameters do not improve the fit quality and lead to instability in eigenvectors, so these are also fixed to the same values as the total sea. Therefore, apart from the normalisation A_+ and the high- x power η_+ , all other parameters in the $s + \bar{s}$ parameterisation are fixed to be the same as the total sea S .

The results of the fit will be discussed in much more detail in Section 8, but at input there is about a 35% normalisation suppression of the strange distribution compared to the average of \bar{u} and \bar{d} at Q_0^2 and some additional high- x suppression. Hence the suppression at the borderline nonperturbative scale $Q_0^2 = 1 \text{ GeV}^2$ is now ~ 0.3 which is the value used in hadronisation models (i.e. the probability to generate $s\bar{s}$ pairs compared to $u\bar{u}$ or $d\bar{d}$ pairs). This is shown in Fig. 20(a) along with the suppression of $c + \bar{c}$ at $Q^2 = 15 \text{ GeV}^2$, i.e. evolved through ~ 7 – 8 times the mass scale. This charm suppression is very similar to that for $s + \bar{s}$ at $Q^2 = 1 \text{ GeV}^2$, implying a suppression due to evolution from slightly more than 0.1 GeV^2 , similar to m_s^2 . We would not of course expect exact correspondence, especially since the strange evolution is in a nonperturbative regime; however, the comparison is surprisingly good except that $c + \bar{c}$ is more suppressed at $x \sim 0.1$ (which is the implication for $s + \bar{s}$ from recent HERMES data on K^\pm production [142]). This gives us confidence that the difference between the total sea and the strange distribution is mainly due to mass suppression, and that our limited parameterisation is well motivated. In Fig. 20(b) we compare the ratio $(s + \bar{s})/(\bar{u} + \bar{d})$ at $Q^2 = 5 \text{ GeV}^2$ obtained from the present analysis to the same quantity from the MRST 2001 [16] and CTEQ6.6 [87] analyses. The MRST 2001 ratio is slightly larger and has much smaller uncertainties, since the input was taken in the fixed form (69) with $\kappa = 0.5$. The CTEQ6.6 ratio in Fig. 20 is much larger for $x \lesssim 0.01$, and has larger uncertainties, due to the more flexible parameterisation in a region where there are no constraining data, but is surprisingly also larger in the region $0.01 \lesssim x \lesssim 0.2$ where the dimuon data should provide a constraint.

The strange asymmetry is not at all well-constrained and the parameterisation,

$$xs_v(x, Q_0^2) \equiv xs(x, Q_0^2) - x\bar{s}(x, Q_0^2) = A_- x^{\delta_-} (1-x)^{\eta_-} (1-x/x_0), \quad (71)$$

is the simplest we can think of with the correct features. The number sum rule of zero strangeness (14) determines x_0 , and hence it is a function of the other parameters. In practice the data are not sufficient to constrain A_- and δ_- independently, so we fix the latter to $\delta_- = 0.2$. Eigenvector 20 (at NLO and NNLO) consists mainly of η_- , hence its peculiar nature shown in Fig. 6. There is little constraint from data, but the allowed values of the parameter are limited.

Summary

In summary, we parameterise with simple forms which do not allow bumps or shoulders in general. Some distributions are restricted by theory assumptions, i.e. the distributions stay extremely small ($\bar{d} - \bar{u}$) or behave like mass-suppressed quarks (strange). We believe that we obtain a good representation of the possible uncertainty in regions where there is a data constraint, since there is an enormous amount of constraining data. We concede that we may underestimate the uncertainty in regions of extrapolation, e.g. the small- x valence quarks, $\bar{d} - \bar{u}$ and strange distributions. However, for some of these the uncertainty is essentially unlimited without making *some* theory assumptions.

6.6 Fit to a reduced dataset

As part of the first proceedings of the HERA–LHC workshop [71], a “benchmark” PDF fit was proposed by S. Alekhin consisting of a limited number of DIS data sets with fairly conservative cuts of $Q^2 \geq 9 \text{ GeV}^2$ and $W^2 \geq 15 \text{ GeV}^2$. The original aim was to provide a test of different fitting codes starting from a common input and theoretical assumptions and fitting to the same data, analogous to the benchmark tables for PDF evolution provided by G. Salam and A. Vogt [70, 71]. The benchmark fit was also carried out by R. Thorne [71] using the MRST fitting code, albeit with some differences in the data sets used and the treatment of correlated systematic errors, but good agreement was found with the benchmark PDFs of S. Alekhin. The benchmark PDFs were also compared to those from the MRST 2001 NLO global fit, and significant discrepancies were found, well outside the error bands, suggesting some inconsistency of the data sets included in the global fit compared to those in the benchmark fit, and an inadequacy of the method for error propagation [71]. However, the benchmark PDF fit of Alekhin/Thorne used a more restrictive parameterisation for the input gluon and sea quark distributions than that used in the global fits. Also, the one-sigma PDF uncertainties for the benchmark fit were defined by $\Delta\chi^2 = 1$, which is clearly inadequate. The benchmark fit was repeated using the MSTW 2008 analysis framework as part of a contribution to the second HERA–LHC proceedings [143]. Although the discrepancy between the benchmark and global PDFs was reduced compared to the previous study [71], there were still sizeable discrepancies in the gluon at low- x and in the down valence distribution. It was realised that some of these discrepancies were due to different values of α_S used and due to the unnecessary assumption that $\bar{d} = \bar{u}$ in the benchmark fit. It is therefore interesting to repeat the benchmark fit here, but fixing α_S to the value in the MSTW 2008

Data set	Reduced fit	Global fit
BCDMS μp F_2 [32]	125 / 157	177 / 157
NMC μp F_2 [33]	72 / 67	75 / 67
NMC $\mu n/\mu p$ [103]	59 / 73	65 / 73
H1 96–97 e^+p NC [109]	46 / 80	47 / 86
ZEUS 96–97 e^+p [112]	155 / 206	163 / 206
All data sets	458 / 589	526 / 589

Table 6: $\chi^2/N_{\text{pts.}}$ for the data sets included in the reduced fit compared to the corresponding values in the global fit.

global fit and allowing for $\bar{d} \neq \bar{u}$. In this way, we can investigate potential data set inconsistency and the (in)adequacy of the error propagation by fitting to a reduced number of data sets, but without any bias due to unnecessarily different input assumptions compared to the global fit. An analogous investigation has recently been carried out by the NNPDF group in Section 5.5 of Ref. [137].

We use the same data sets as specified in the MSTW version of the HERA–LHC benchmark fit [143], that is, we fit BCDMS data on F_2^p [32], NMC data on F_2^p [33], NMC data on F_2^d/F_2^p [103], H1 96–97 e^+p NC data on $\tilde{\sigma}$ [109] and ZEUS 96–97 e^+p NC data on $\tilde{\sigma}$ [112], with cuts of $Q^2 \geq 9$ GeV² and $W^2 \geq 15$ GeV². The heavy flavour treatment is the same as in the global fit, that is, using the GM-VFNS with $m_c = 1.4$ GeV and $m_b = 4.75$ GeV. The strong coupling is fixed at the value obtained in the NLO global fit. The input PDF parameterisation is taken to be the same as in the global fit, with the following exceptions. None of the data sets included constrain strangeness, and indeed, negative strange quark distributions are obtained if allowed to go free. Therefore, we fix the input $s_v = s - \bar{s}$ and $(s + \bar{s})/(\bar{u} + \bar{d})$ to be the same as in the global fit.

In Table 6 we compare the χ^2 values for the reduced fit to the χ^2 values for the same data points in the global fit. The deterioration in the quality of the fit (total $\Delta\chi^2 = 68 \gg 1$) with the inclusion of more data sets is evident, especially for the BCDMS F_2^p data. Both the reduced fit and the global fit give an acceptable description of all data sets included in the reduced fit, and so any error propagation based on $\Delta\chi^2 = 1$ is clearly inadequate.

Application of the dynamic tolerance procedure described above to the reduced fit gives values shown in Fig. 21. Typical values for the tolerance are $T \sim 3$ for 68% C.L. uncertainties and $T \sim 5$ for 90% C.L. uncertainties, that is, slightly smaller than the typical values in the global fit, and with less variation between different eigenvector directions. We note that the departure from ideal quadratic behaviour ($T = t$) for higher eigenvector numbers is more severe for the reduced fit than the global fit, since the PDF parameters are much less constrained in the reduced fit and hence there are more relatively “flat” directions in eigenvector space.

In Fig. 22 we compare the u_v and d_v distributions at $Q^2 = 20$ GeV² from the reduced and global fits. The dashed lines indicate the uncertainty bands obtained with $T = 1$, while the outer uncertainty bands are obtained using the dynamic tolerance shown in Fig. 21. For the valence quarks one can see that there are differences in detail, particularly for d_v . Using uncertainties determined by $\Delta\chi^2 = 1$ there are still significant discrepancies in places, e.g. d_v at $x = 0.05$,

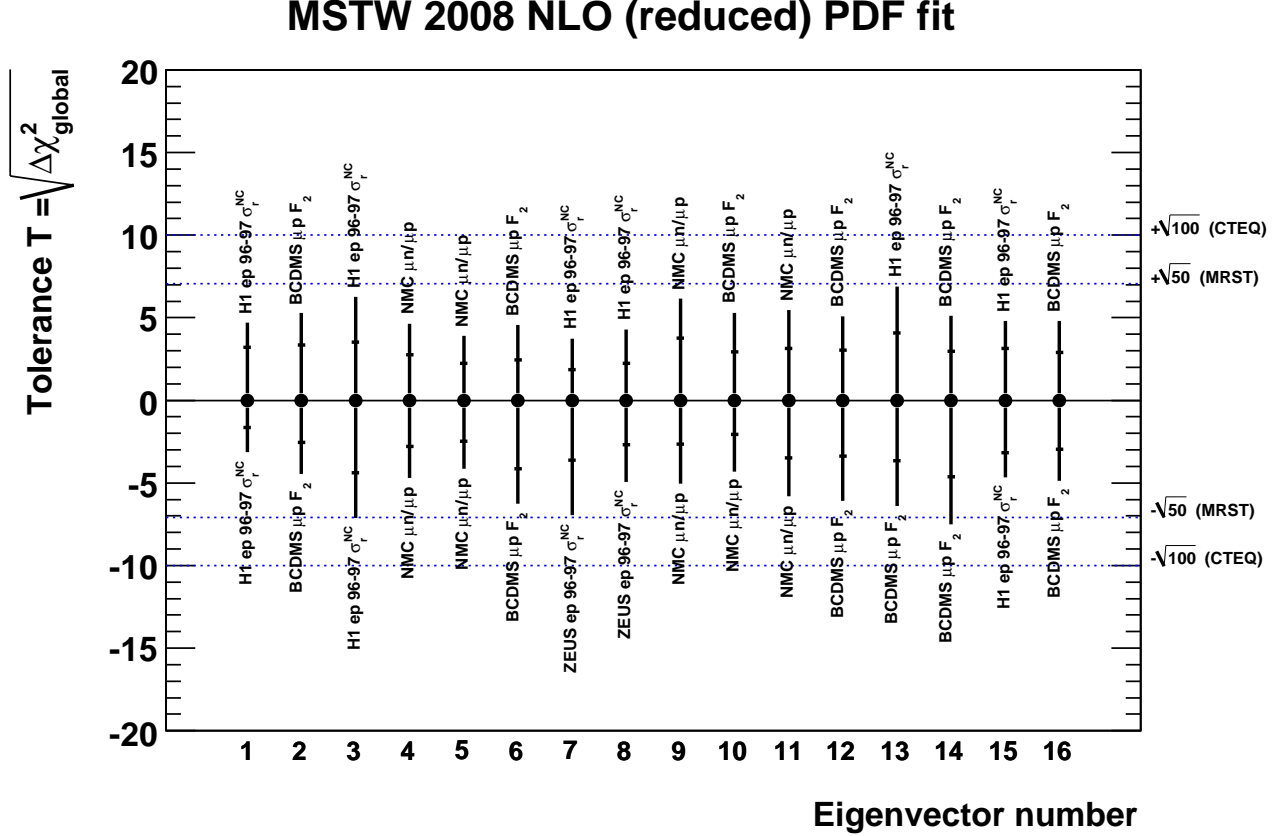
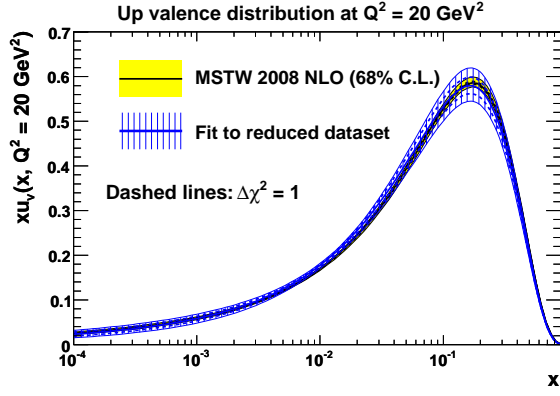


Figure 21: Tolerance for each eigenvector in the reduced fit.

but using the dynamic tolerance the two fits are consistent, though the global fit leads to much smaller uncertainties, as expected. In Fig. 23(a) we see the curves for the \bar{u} distribution, where again there is good compatibility, though the lack of Drell–Yan and neutrino DIS data in the benchmark fit leads to a huge uncertainty at high x . At $x > 0.5$ there is some discrepancy. This is related to the situation for the gluon, seen in Fig. 23(b). At small x the two gluons agree well, again with that from the reduced fit having larger uncertainty. At high x there is some disagreement. Using $\alpha_S(M_Z^2) \simeq 0.120$ the extremely good fit to high- x BCDMS data, seen in Table 6, requires a large high- x gluon in order to obtain the sufficiently flat shape in Q^2 preferred by the data. This produces a pull in clear contradiction to other data, and which is overwhelmed in a global fit, but survives in the reduced fit.

Hence, we can see that the “benchmark” type of fit to a reduced number of data sets proposed in Ref. [71] is a useful means of comparing different fitting approaches, as demonstrated in Refs. [71, 143]. However, the central values of the PDFs obtained and the uncertainties obtained using the Hessian method with the textbook $\Delta\chi^2 = 1$ should not be taken seriously — hence our decision not to make these sets publicly available. As shown above, comparison of the benchmark PDFs with those obtained from more complete fits can be useful in suggesting the reliability

(a)



(b)

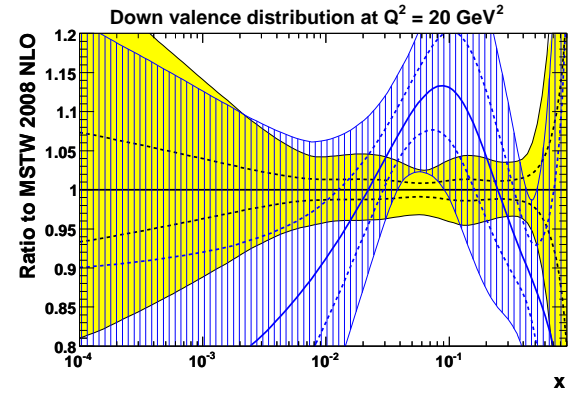
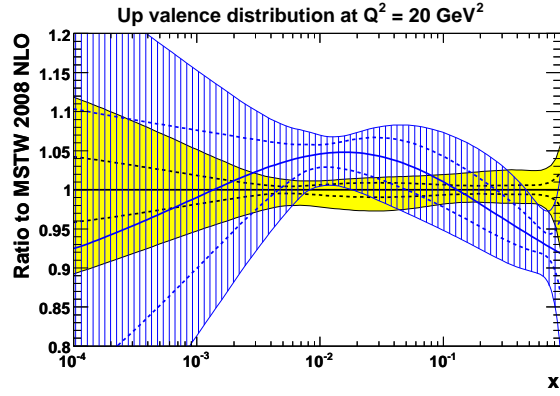
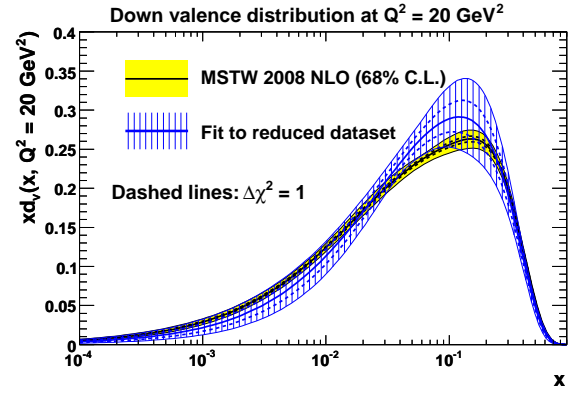



Figure 22: Comparison of the (a) up and (b) down valence quark distributions from the global and reduced NLO fits.

(a)



(b)

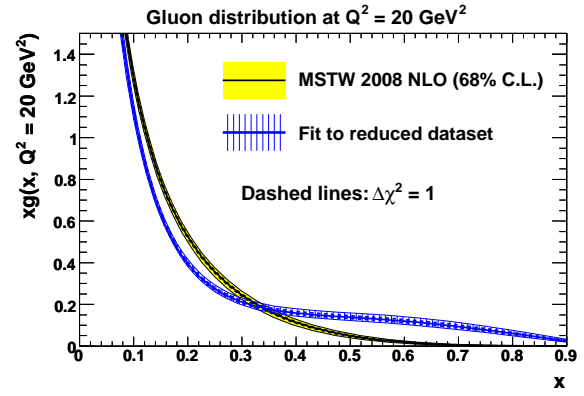
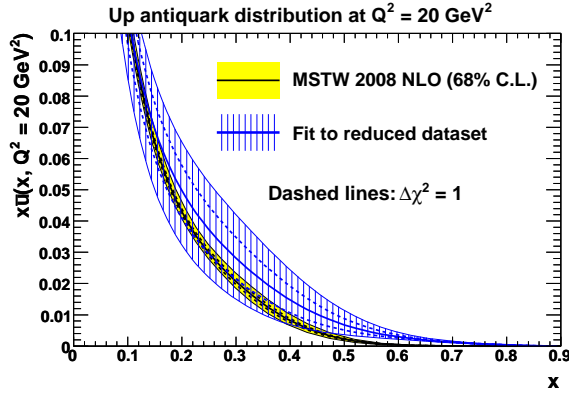
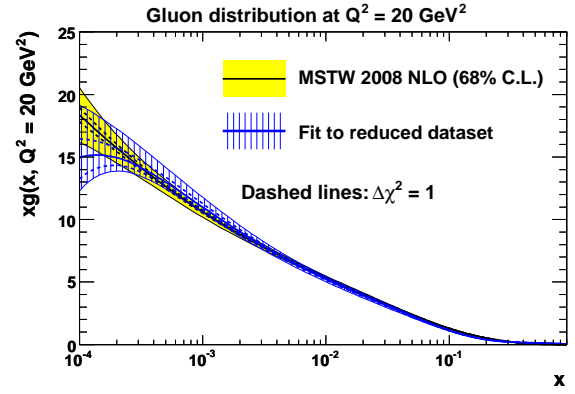


Figure 23: Comparison of (a) the \bar{u} distribution and (b) the gluon distribution from the global and reduced NLO fits.

of more elaborate methods of determining the PDF uncertainty, giving good support for the “dynamic tolerance” approach introduced in this paper. Benchmark fits can also be useful for highlighting data sets which are inconsistent with the rest of the data in a global fit. Indeed, it is shown that there is clear tension between the BCDMS F_2^p data and the rest of the data, which is a concern, although the degree of compatibility becomes better when the cut of $Q^2 \geq 9$ GeV^2 is lowered to our standard choice of $Q^2 \geq 2$ GeV^2 .

7 Description of inclusive structure function data

7.1 HERA structure functions

We show the fit to a selection of the HERA neutral-current data for the reduced cross section,

$$\tilde{\sigma}(x, Q^2) = F_2(x, Q^2) - \frac{y^2}{1 + (1 - y^2)} F_L(x, Q^2), \quad (72)$$

in Fig. 24. This illustrates the general features, and the quality of the fit can be judged from the χ^2 values in Table 2 in Section 5. At each order the quality of the fit is very similar to previous MRST fits. A number of general features may be observed by examination of Fig. 24. It is clear that the LO fit is not able to generate sufficient evolution to give a genuinely good fit, being flatter at both low and high Q^2 in each x bin. This is due to the presence of large positive contributions to the splitting function P_{qq} beyond LO, and even the very large value of α_S in the LO fit cannot fully compensate. We also note that the turnover of $\tilde{\sigma}(x, Q^2)$ at high Q^2 (equivalent to high y) in each x bin is largest at LO, particularly at small x . As we will see later in Section 13 this is because $F_L(x, Q^2)$ is very large at LO. The NLO and NNLO fits are very similar to each other. The main difference is that the NNLO curves are rather steeper at low x and Q^2 , being driven by an extra $\ln(1/x)$ factor in the NNLO P_{qq} . Although it is difficult to see from the plot this does actually lead to a worse fit at NNLO for the H1 97 minimum bias e^+p data [109], see Table 2, the data preferring a slightly flatter shape. Overall, the quality of the fit is very good. The only flaw is a slight trend for the evolution to not persist quite strongly enough at high Q^2 for $x \sim 0.001$ – 0.01 , although there is quite a large spread between the H1 and ZEUS data at high Q^2 in these x bins. This might possibly be a sign of further required corrections to the theory of some sort, and indeed an improvement due to small- x resummation was shown to be possible in Ref. [144].

The only HERA structure function data completely new to our analysis are the H1 MB 99 e^+p data [31] — a handful of points which probe the highest y values and hence display sensitivity to $F_L(x, Q^2)$. In Fig. 25 we show our results in bins of Q^2 to best compare to these specific data. The dashed lines show the one-sigma PDF uncertainty bands. We noted in Ref. [28] that while the turnover at small x (high y) did not occur with our central NLO fit, it did occur with the central NNLO fit due to the increase of $F_L(x, Q^2)$ from the NNLO coefficient function. It is striking that the central value of the NNLO result no longer fully turns over at low Q^2 and small x . However, from Table 2 one can see that the fit to the H1 MB 99 e^+p data is still a little better at NNLO, but also that it is not very bad at LO or NLO. The data, while being indicative, are not in practice very constraining. It is also clear from the plot that the full turnover is allowed within the uncertainties of the PDFs. We will discuss more direct results about, and relating to, $F_L(x, Q^2)$ in Section 13.

Charged-current data from e^+p scattering [35, 36] are also included in the fit. They are fit perfectly well, as seen in Table 2, but at present do not provide a strong constraint on the fit. Much more precise data are expected from HERA II analyses.

MSTW 2008




Figure 24: The reduced cross section $\tilde{\sigma}(x, Q^2)$ for a selection of HERA data and the structure function $F_2(x, Q^2)$ for a selection of fixed-target data. The curves represent $\tilde{\sigma}(x, Q^2)$, evaluated at the HERA centre-of-mass energy, and so for the high x and low Q^2 values corresponding to fixed-target data, y is negligible and the curves effectively represent the required $F_2(x, Q^2)$. The normalisations of data sets are those in the best fit at NNLO, found in Table 3. These are very similar to those at NLO, but at LO the HERA data would be 0.5–1% lower in general.




Figure 25: Predictions for the reduced cross section compared to HERA data in bins of Q^2 . The dashed lines show the one-sigma PDF uncertainty bands.

7.2 Fixed-target neutral-current structure functions

The fit to a selection of fixed-target data is also shown in Fig. 24, where in this case the data and the curve (to an extremely good approximation) represent the structure function $F_2(x, Q^2)$. These data have appeared in numerous previous MRST analyses, and at NLO and NNLO the fit is of the same quality as previously. These data are one of the main constraints on the value of $\alpha_S(M_Z^2)$, though it is correlated to the high- x gluon. Indeed, the value of $\alpha_S(M_Z^2)$ at NNLO is slightly smaller than in our most recent analyses [21], and so the high- x gluon is correspondingly smaller. Perhaps the most striking new feature is that at LO the value of α_S has increased sufficiently that the fit quality is much closer to that at NLO and NNLO, the size of the coupling largely making up for the increased evolution at higher orders in the high- x regime. Indeed, at $x = 0.35$ we see that at very low Q^2 (below our W^2 cut on data) the LO curve in Fig. 24 is falling most quickly — a feature not seen previously.

7.3 Neutrino structure functions

There have been significant changes in both the data to which we fit, and the details of the procedure used, for neutrino structure functions. The NuTeV data [37] on $F_2 \equiv F_2^{\nu N} = F_2^{\bar{\nu} N}$ and $xF_3 \equiv (xF_3^{\nu N} + xF_3^{\bar{\nu} N})/2$, where $N = (n + p)/2$, replace the previous CCFR data [145]. These two experiments both use an iron target, corrected to an isoscalar target, and cover much the same kinematic range, but the NuTeV data are more precise. The two sets agree fairly well, except in the very high- x (i.e. $x > 0.5$) region, where the main source of the discrepancy is the calibration of the magnetic field map of the muon spectrometer. In practice we find the high- x NuTeV data very difficult to fit. At high- x the predictions are mainly determined by the valence up quark distribution, which is very well constrained by the neutral-current structure function data. Given the degree of experimental uncertainty in this region, we do not fit to neutrino data with $x > 0.5$. The main information from these data comes from $x \lesssim 0.3$, where we are sensitive to the details of the different valence and sea quark contributions, which are weighted differently for neutrino structure function data compared to neutral-current data. We also include the recent CHORUS F_2 and xF_3 data [38] which cover a similar range in x but at slightly lower Q^2 on average, and are taken using a lead target. These are completely compatible with the NuTeV data for $x < 0.5$, but seem to lie a little lower at the highest common x value of 0.65. We apply the same cut to the CHORUS data of $x > 0.5$ for consistency. As mentioned in Section 5 we apply a cut of $W^2 \geq 25 \text{ GeV}^2$ on the xF_3 data to guard against potentially larger (than F_2) higher-twist contributions.

In the analysis of these data, we have adopted a more sophisticated treatment of nuclear corrections compared to previous fits. We apply the nuclear corrections R_f , defined as $f^A(x, Q^2) = R_f(x, Q^2, A)f(x, Q^2)$, separately for each parton flavour f using the results of a NLO fit by de Florian and Sassot (nDS) [147], using the GRV98 [148] free proton PDFs as input; see Fig. 26. Here, f^A are defined to be the parton distributions of a *proton* bound in a nucleus of mass number A . In the figure we compare to a NLO analysis by Hirai, Kumano and Nagai (HKN07) [146] using the MRST98 [12] free proton PDFs as input, and to the simple “RATFE” flavour- and Q^2 -independent corrections [12], applied in previous MRST fits. Since

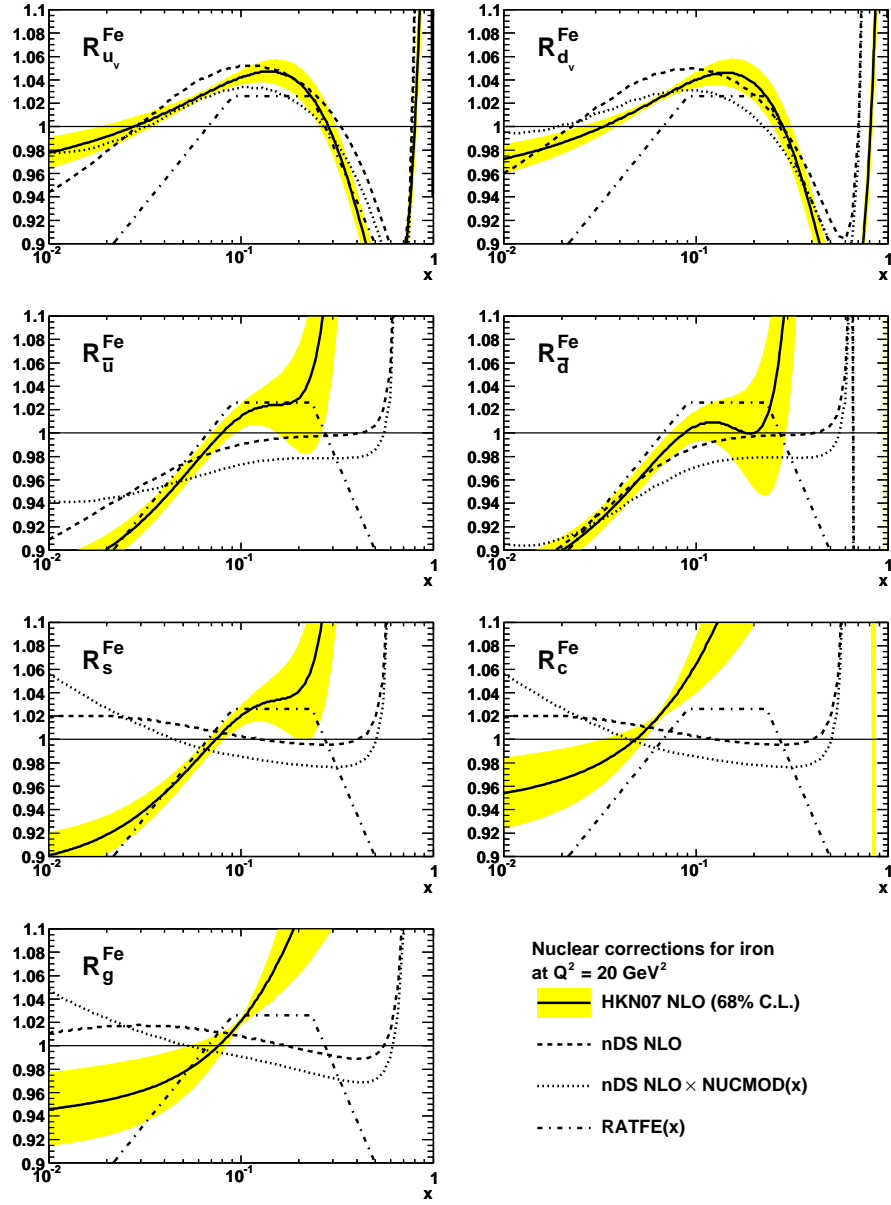


Figure 26: Nuclear corrections for iron ($A = 56$, $Z = 26$), at $Q^2 = 20 \text{ GeV}^2$. The solid lines, and the associated one-sigma uncertainty bands, are from an analysis by Hirai, Kumano and Nagai (HKN07) [146]. The dashed lines are from an analysis by de Florian and Sassot (nDS) [147], and the dotted lines are multiplied by the correction factor, (73), determined from our NLO global fit. We take the nDS nuclear corrections for charm (and bottom) quarks to be the same as those for strange quarks. The flavour- and Q^2 -independent “RATFE(x)” corrections [12], shown by the dot-dashed lines, were applied in previous MRST fits.

the nDS analysis [147] neglects heavy flavours, we take the nuclear corrections for charm (and bottom) quarks to be the same as those for strange quarks. The nDS and HKN07 nuclear corrections are similar for the up and down quarks, which are most important for the neutrino structure functions, but differ significantly for the less important gluon and other quark flavours.²⁰ An alternative approach consists of extracting nuclear PDFs from data on a single nuclear target (iron) [150].

The nDS nuclear corrections do not have an associated uncertainty. We assume that the uncertainty is likely to be of order a few percent, and account for it by multiplying the nuclear corrections $R_f(x, Q^2, A)$ by a flavour-independent modification function of the form:

$$\text{NUCMOD}(x) = \begin{cases} (1 + 0.03 r_1)[1 + 0.015 r_2 \ln^2(x_m/x)] & : x < x_m \\ (1 + 0.03 r_1)[1 + 0.015 r_3 \ln^2(x/x_m)] & : x \geq x_m \end{cases}, \quad (73)$$

where $x_m = \exp(-2.5)$ is chosen to be roughly in the middle (logarithmically) of the x range spanned by the neutrino data. The parameters r_1 , r_2 and r_3 are allowed to go free when we obtain our fits. This function allows, for example, a little more variation than the uncertainties in the correction factors quoted in HKN07 [146]. The extent to which the central values of the nDS nuclear corrections are modified by (73) is shown in Fig. 26: 2–3% is typical, and we regard this as perfectly acceptable.

The quality of the NNLO fit to F_2 data is shown in Fig. 27, where the data which fail to satisfy the W^2 cut and the $x < 0.5$ cut are also shown. The fit quality is perfectly good, as seen in Table 2, and the NuTeV and CHORUS data agree very well in the region of overlap, the latter lying a little lower in Q^2 . For $x > 0.5$ there is a slight tendency for CHORUS data to lie above the prediction, which is dominated by the u_v quark determined by neutral-current DIS data. The NuTeV data, however, lie consistently well above the prediction.

The quality of the NNLO and LO fits to xF_3 data is shown in Figs. 28 and 29, respectively. One obvious result of the LO fit is that the Q^2 slope of the theory is consistently too flat, except for $x > 0.275$, and the fit is correspondingly poor. This failure is largely cured in the NNLO fit, though perhaps not completely at the lowest x values. Indeed, it is clear from the χ^2 values in Table 2 that the fit to xF_3 data, for both experiments, improves significantly from LO to NLO, and then even more at NNLO. Since

$$xF_3 \equiv \frac{1}{2} (xF_3^{\nu N} + xF_3^{\bar{\nu} N}) = C_3 \otimes \sum_i (q_i - \bar{q}_i), \quad (74)$$

this is a particularly clean test of perturbative QCD (and α_S) with no complications from mixing in parton evolution — once the single parton combination is determined at one scale we have a complete prediction (in terms of α_S) at higher Q^2 . In this case the improvements due to the change in shape with Q^2 of the structure function at increasing order are clearly seen. As seen from Table 2, the χ^2 values for the 45 NuTeV xF_3 points improve from $62 \rightarrow 40 \rightarrow 34$ in going from LO \rightarrow NLO \rightarrow NNLO, while the χ^2 values for the 33 CHORUS xF_3 points improve from $44 \rightarrow 31 \rightarrow 26$. It is not certain that the improvement has converged at NNLO, and

²⁰The EPS09 [149] corrections are in good agreement with those from nDS [147] for valence and sea quarks.

Ratio of NuTeV and CHORUS F_2 data to MSTW 2008 NNLO PDF fit

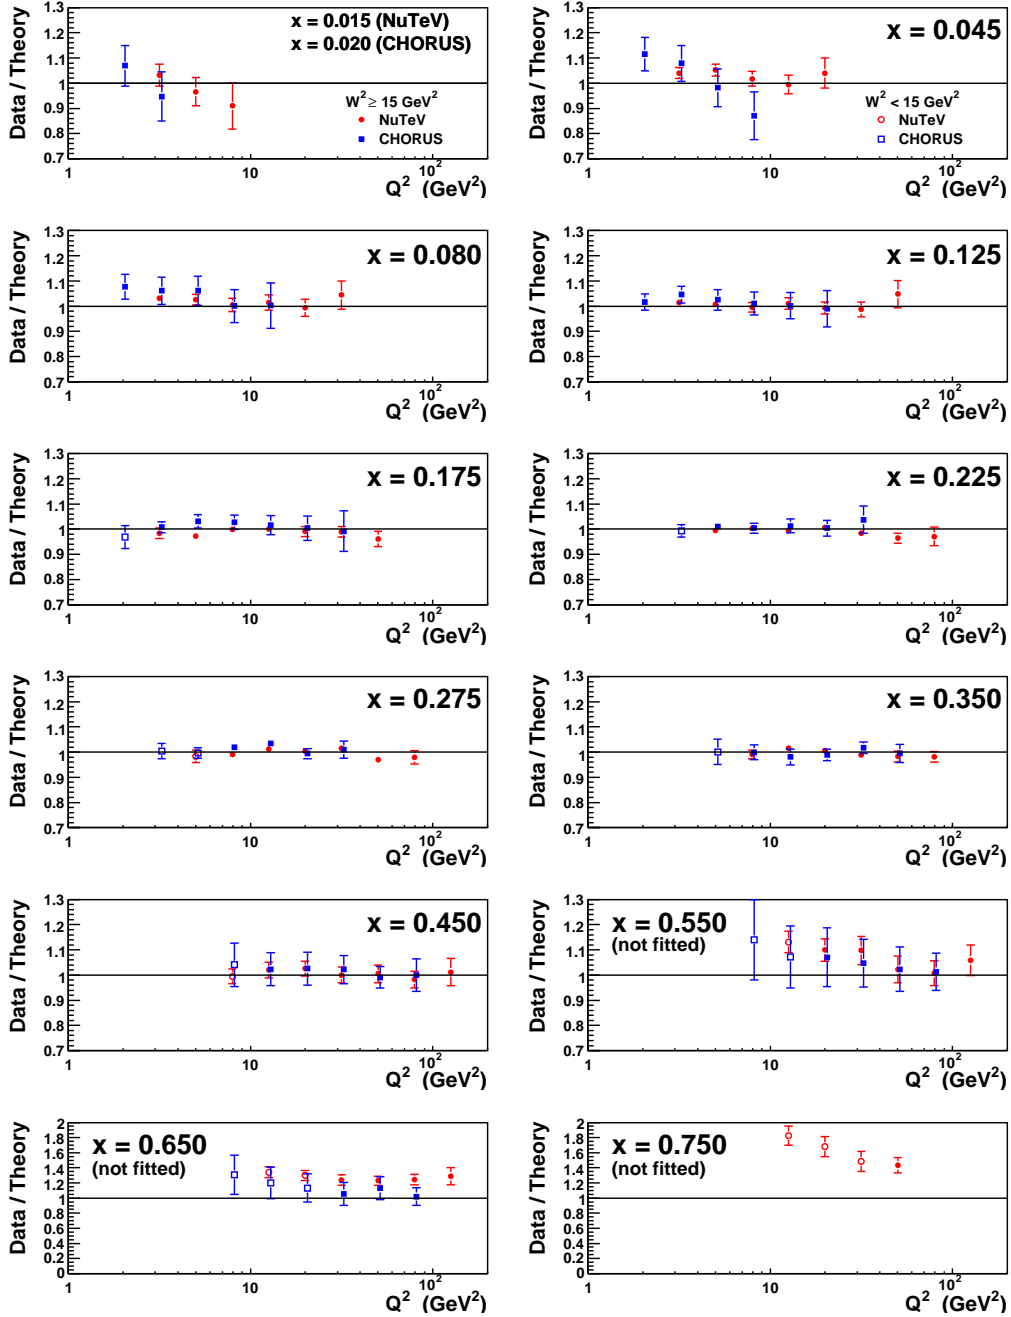


Figure 27: The quality of the NNLO fit to $F_2 \equiv F_2^{\nu N} = F_2^{\bar{\nu} N}$ data from NuTeV [37] and CHORUS [38]. The data which fail to satisfy the W^2 cut and the $x < 0.5$ cut are also shown.

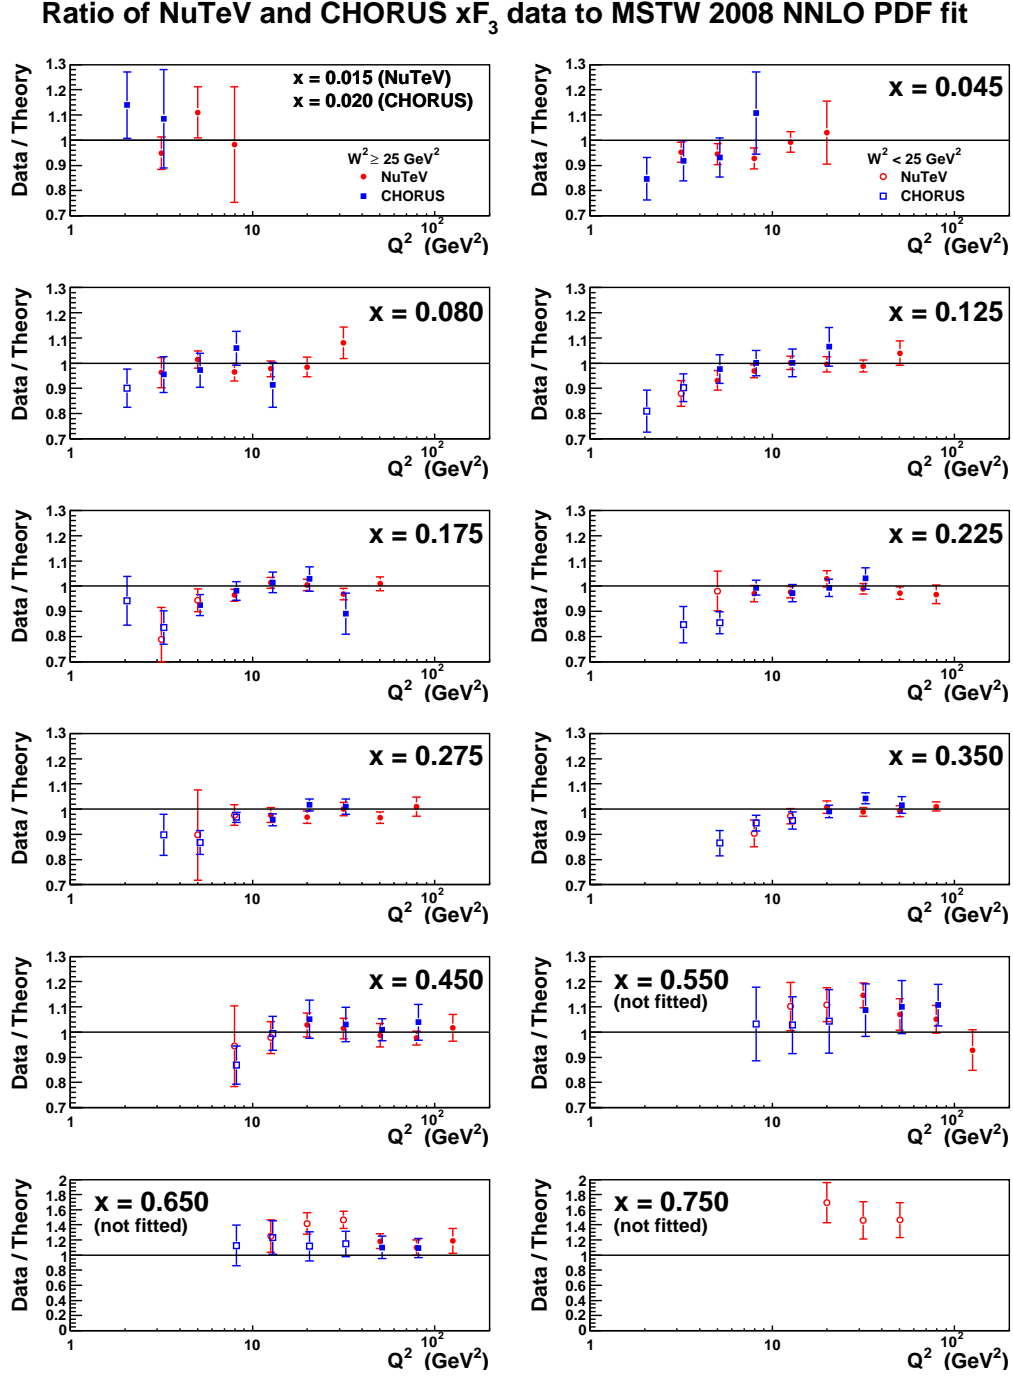


Figure 28: The quality of the NNLO fit to $xF_3 \equiv (xF_3^{\nu N} + xF_3^{\bar{\nu} N})/2$ data from NuTeV [37] and CHORUS [38]. The data which fail to satisfy the W^2 cut and the $x < 0.5$ cut are also shown.

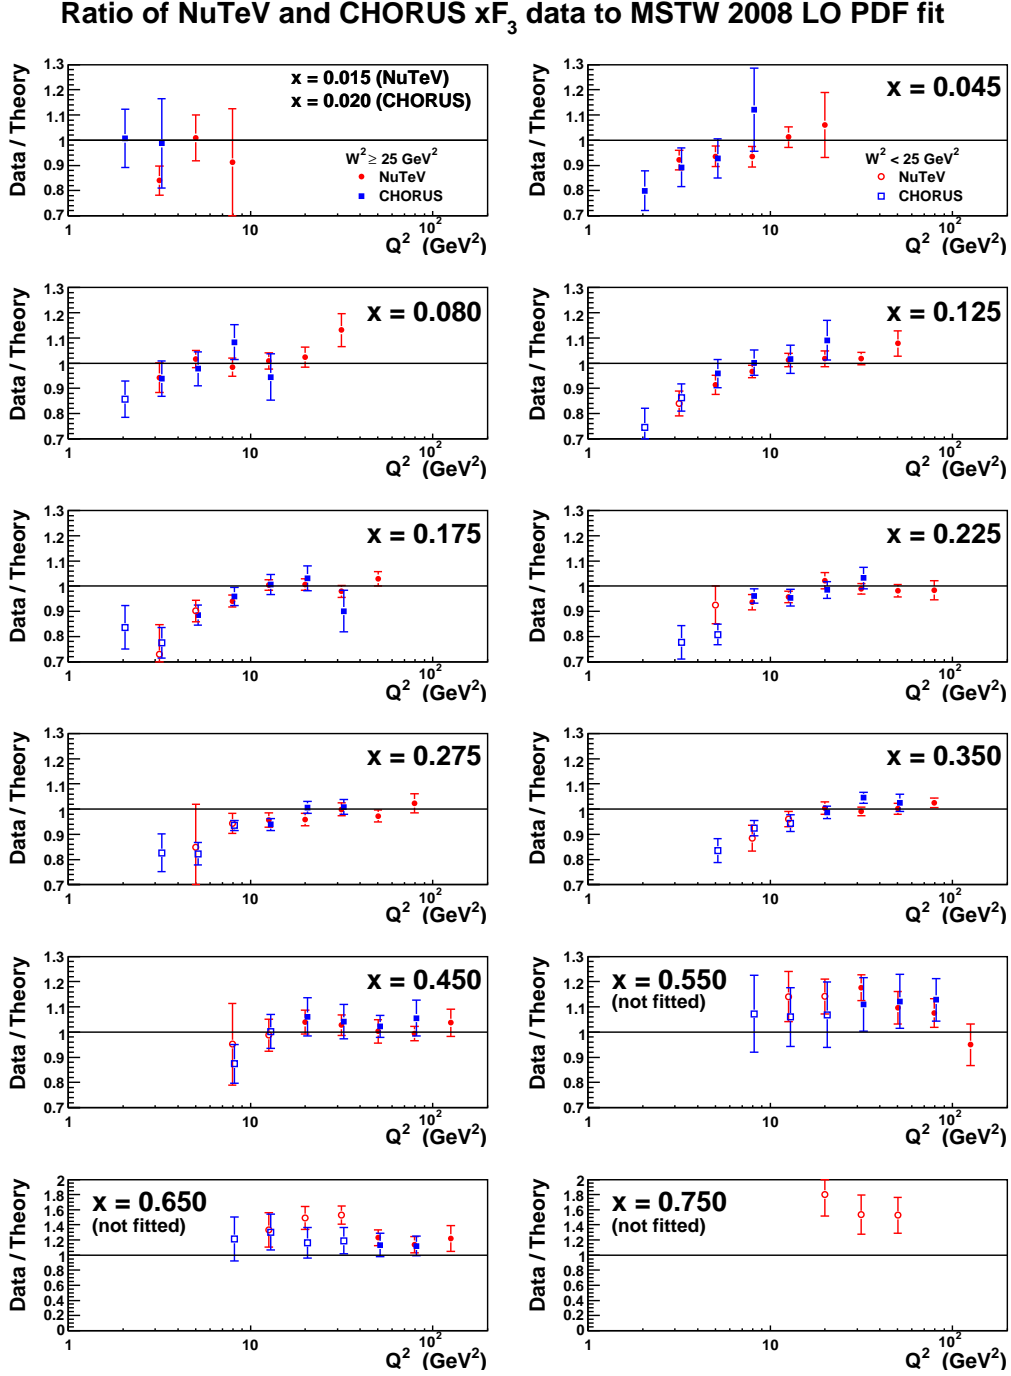


Figure 29: The quality of the LO fit to $xF_3 \equiv (xF_3^{\nu N} + xF_3^{\bar{\nu} N})/2$ data from NuTeV [37] and CHORUS [38]. The data which fail to satisfy the W^2 cut and the $x < 0.5$ cut are also shown.

there is perhaps a little scope for higher-order corrections to improve matters further. This is quite possibly the case for many quantities, but the mixing between parton flavours obscures the interpretation for most. It is also interesting that these data favour fairly high values of $\alpha_S(M_Z^2)$, contrary to the claim often made about the results obtained from fitting structure function data alone [139, 151, 152]. The trend of the xF_3 data at $x > 0.5$ to lie above the theory predictions is the same as for F_2 . In Figs. 28 and 29 we also show the data below our higher (than for F_2) $W^2 \geq 25 \text{ GeV}^2$ cut. We see that for $x < 0.35$ these data do indeed lie systematically below the NNLO result, implying a large negative higher-twist contribution, which may also very slightly influence data above our cut. We leave a systematic analysis of such higher-twist contributions to a future study.

8 Neutrino dimuon production: constraints on strangeness

Our new analysis allows much more flexibility for the strange quark and antistrange quark distributions. This is both for the total and the two separate contributions. Protons have no valence strange quarks, that is,

$$\int_0^1 dx [s(x, Q^2) - \bar{s}(x, Q^2)] = 0. \quad (75)$$

But this does *not* necessarily mean that $s(x, Q^2) = \bar{s}(x, Q^2)$ for all x . For example, in a meson cloud model, the proton can fluctuate into a $\Lambda(uds)$ and a K or $K^*(u\bar{s})$, where the \bar{s} in the K or K^* carries a larger momentum fraction of its parent hadron than the s quark in the Λ . Nevertheless, the assumption in earlier PDF fits has been that the input distributions satisfy

$$s = \bar{s} = \frac{\kappa}{2}(\bar{u} + \bar{d}), \quad (76)$$

with a constant $\kappa \approx 0.4\text{--}0.5$, justified by a simplified analysis of dimuon data by the CCFR experiment [40]. Updated CCFR and NuTeV dimuon cross sections are now available, therefore we include these data in the global fit to constrain s and \bar{s} separately; see, for example, the process $\nu_\mu N \rightarrow \mu^+ \mu^- X$ proceeding via the subprocess $W^+ s \rightarrow c$ shown in Fig. 30. Previous studies have been made in Refs. [153–156].




Figure 30: Dimuon production in neutrino–nucleon DIS from charm production via scattering of the W^+ boson off a strange or down quark from a nucleon in the heavy nuclear target N . Diagram taken from Ref. [39].

8.1 Description of dimuon cross sections

The CCFR and NuTeV dimuon cross sections, $\nu_\mu N \rightarrow \mu^+ \mu^- X$, can be related to the underlying process, $\nu_\mu N \rightarrow \mu^- c X$, using

$$\frac{d^2\sigma}{dx dy}(\nu_\mu N \rightarrow \mu^+ \mu^- X) = B_c \mathcal{A} \frac{d^2\sigma}{dx dy}(\nu_\mu N \rightarrow \mu^- c X), \quad (77)$$

and similarly for the antineutrino data. Here, $B_c = 0.099 \pm 0.012$ is the charm semileptonic branching ratio [153]. The kinematic acceptance correction \mathcal{A} accounts for the 5 GeV cut on the energy of the charm decay muon and was calculated [153] for each bin using the DISCO Monte Carlo simulation [157] assuming a charm mass of 1.4 GeV and Collins–Spiller [158] fragmentation parameter $\epsilon = 0.6$. Allowing for the CKM mixing, the cross section has the $\mathcal{O}(\alpha_S^0)$ form:

$$\frac{d\sigma}{dx dy}(\nu_\mu N \rightarrow \mu^- c X) \propto |V_{cs}|^2 \xi s^A(\xi, Q^2) + |V_{cd}|^2 \xi \left[\frac{Z}{A} d^A(\xi, Q^2) + \left(1 - \frac{Z}{A}\right) u^A(\xi, Q^2) \right], \quad (78)$$

where the rescaling variable $\xi \equiv x(1 + m_c^2/Q^2)$, and where $Z = 23.403$ and $A = 49.618$ for the CCFR/NuTeV target [153]. Thus, the ν_μ and $\bar{\nu}_\mu$ dimuon cross sections are direct measures of s and \bar{s} , respectively. We have assumed isospin symmetry, $d^n = u^p \equiv u$, and, as usual, have taken the renormalisation and factorisation scales to be $\mu_R = \mu_F = Q$. Finally, just as for the inclusive neutrino data, we apply the nuclear corrections of de Florian and Sassot (nDS) [147] for each parton flavour, and hence relate the nuclear PDFs, $f^A(x, Q^2)$, to the proton PDFs, $f(x, Q^2)$; see Section 7.3. As evident from the difference between the nDS [147] and HKN07 [146] parameterisations for nuclear corrections to strange quarks shown in Fig. 26, the uncertainty is likely to be a few percent. However, this uncertainty is smaller than the experimental uncertainty on $s + \bar{s}$ at the relatively low scales where the data are probed. Additionally, there is no evidence for any systematic problem with shape in the fit quality.

It is convenient to parameterise $s \pm \bar{s}$, rather than s and \bar{s} separately, as explained in Section 6. Recall that at the input scale $Q_0^2 = 1 \text{ GeV}^2$ we take

$$x(s + \bar{s}) = A_+ x^{\delta_S} (1 - x)^{\eta_+} (1 + \epsilon_S \sqrt{x} + \gamma_S x), \quad (79)$$

$$x(s - \bar{s}) = A_- x^{\delta_-} (1 - x)^{\eta_-} (1 - x/x_0), \quad (80)$$

where the parameters δ_S , ϵ_S and γ_S are in common with the parameterisation of the total light quark sea, S , see (8). The final factor in (80) is to ensure that the number sum rule (75) is satisfied, and implies that $(s - \bar{s})$ changes sign at $x = x_0$.

The description of the CCFR and NuTeV neutrino and antineutrino dimuon data given by the NNLO fit is shown in Figs. 31, 32, 33 and 34. Clearly the quality of the fit is very good. There is perhaps a slight tendency for the CCFR antineutrino data to lie a little above the curves and the NuTeV antineutrino data to lie a little below the curves, suggesting that the more recent NuTeV data favour a lower $\bar{s}(x, Q^2)$ distribution in the region of the data and are mainly responsible for the asymmetry. The results are similar at LO, NLO and NNLO, as suggested by the quality of the fits shown in Table 2. Note that the branching ratio $B_c = 0.099$ is fixed in the fits. If allowed to go free there is a distinct tendency for it to choose a high value. This is the same sort

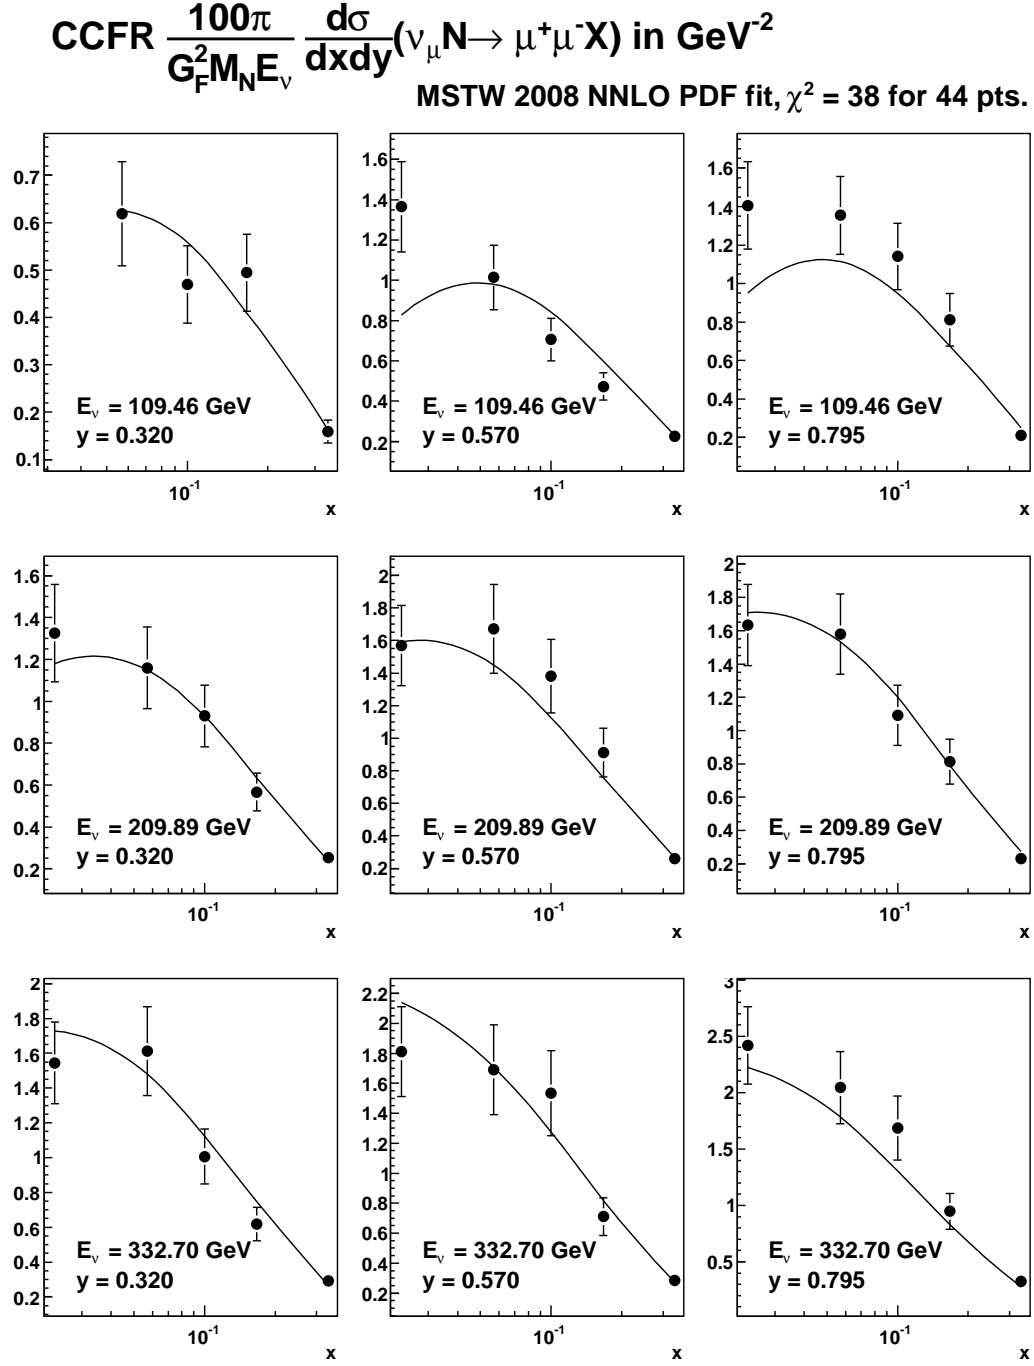


Figure 31: The quality of the NNLO fit to the CCFR neutrino-initiated dimuon production.

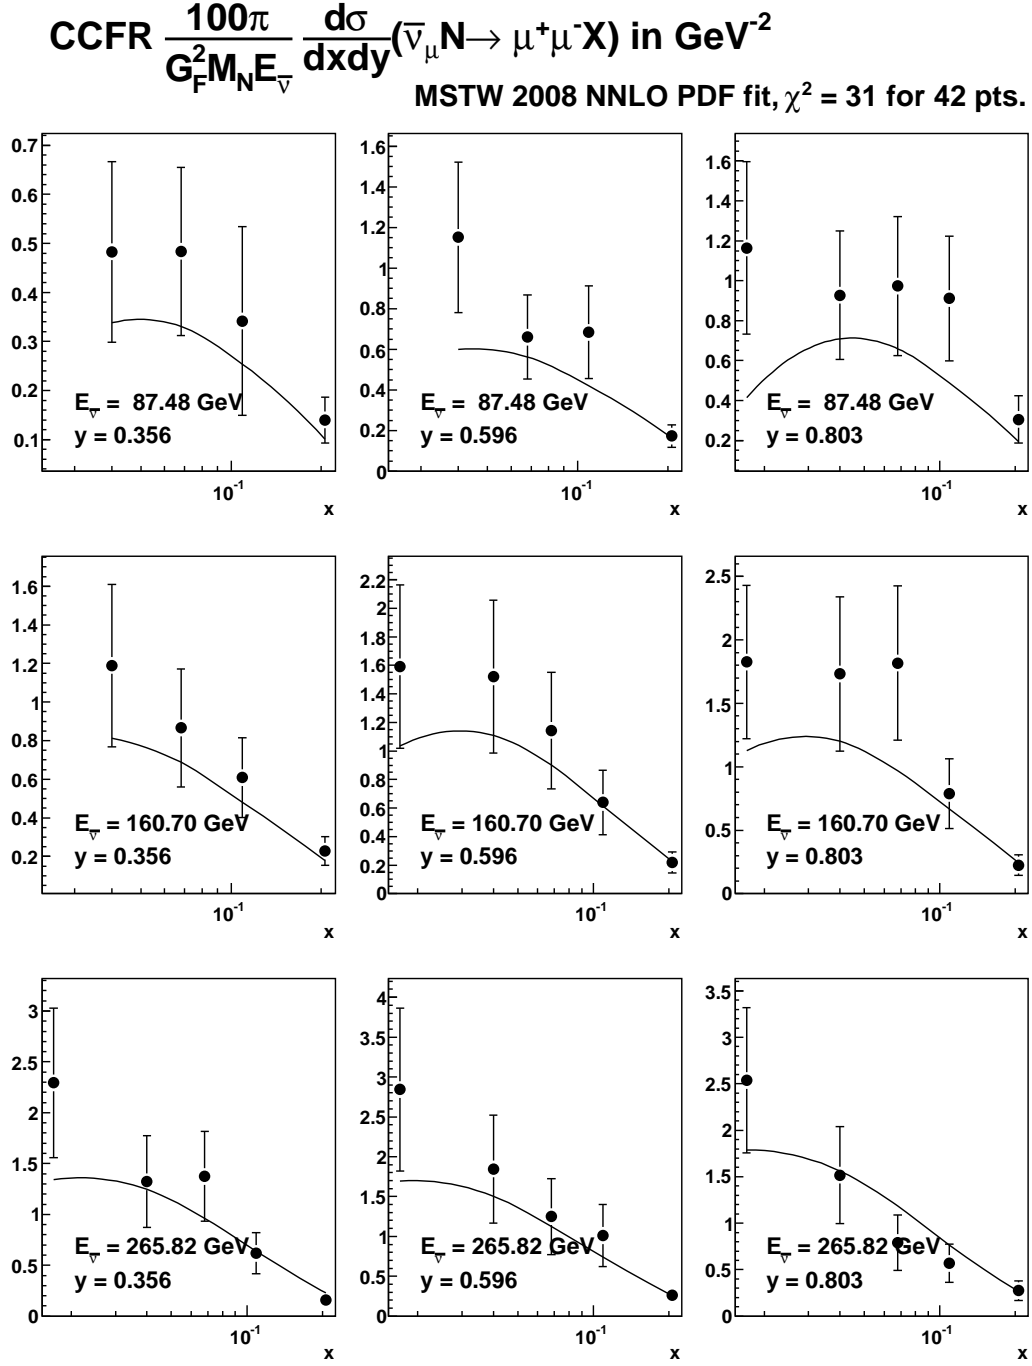


Figure 32: The quality of the NNLO fit to the CCFR antineutrino-initiated dimuon production.

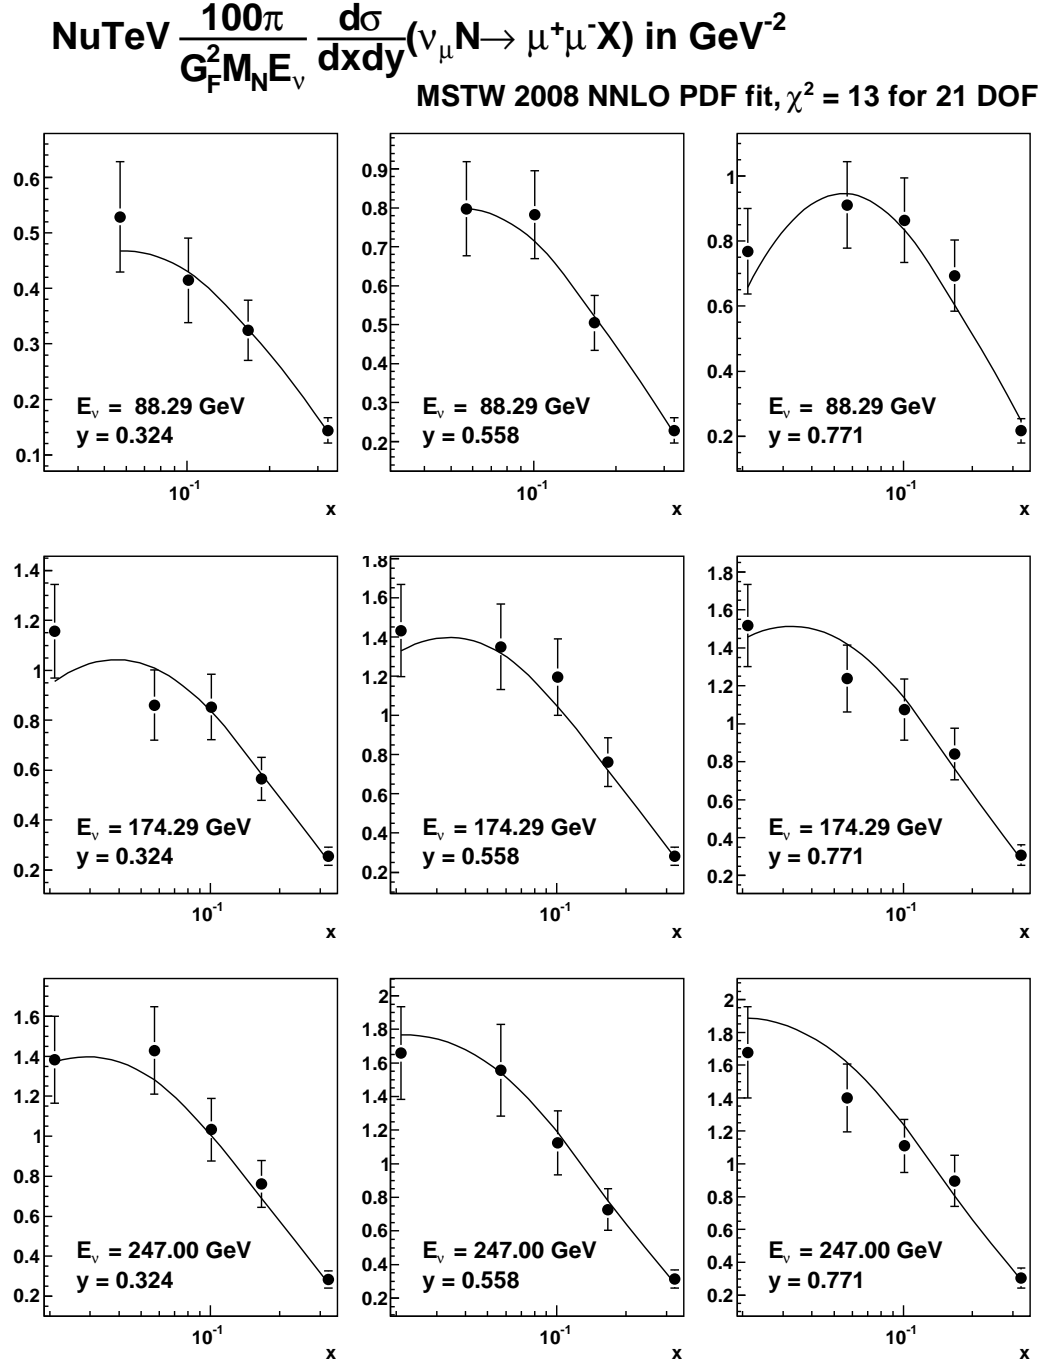


Figure 33: The quality of the NNLO fit to the NuTeV neutrino-initiated dimuon production.

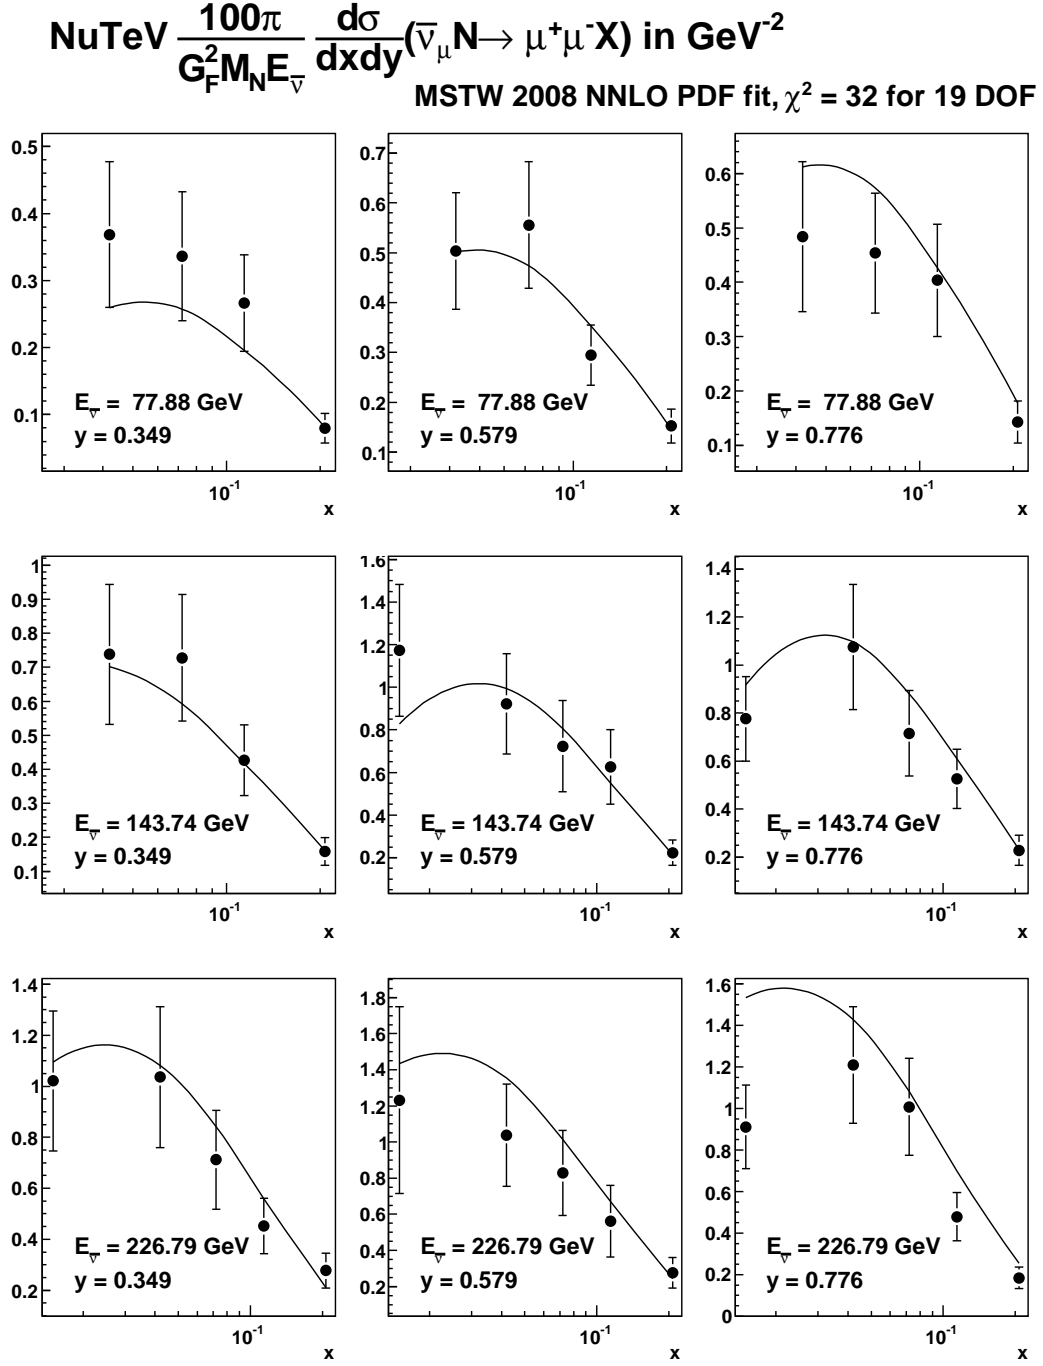


Figure 34: The quality of the NNLO fit to the NuTeV antineutrino-initiated dimuon production.

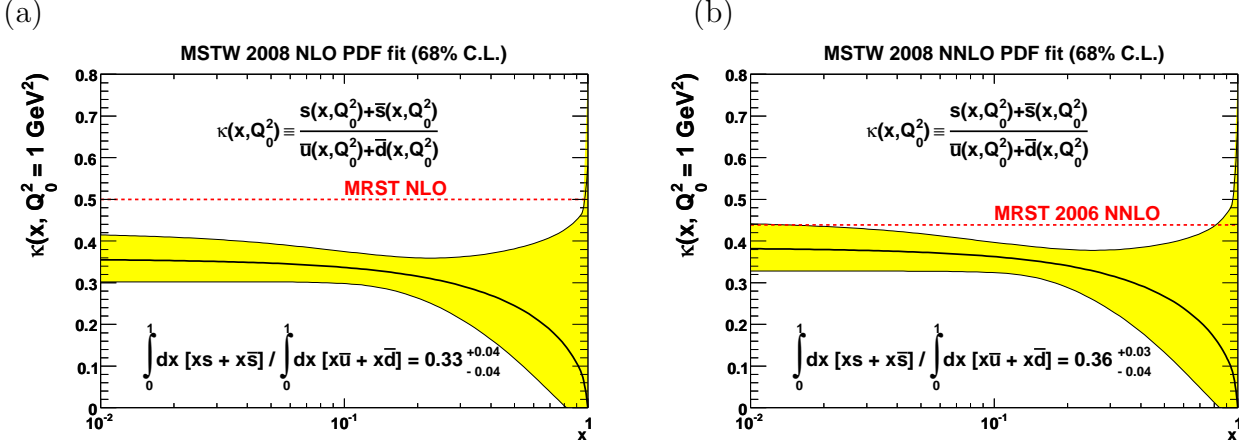


Figure 35: Ratio of the strange to the non-strange distributions at the input scale $Q_0^2 = 1 \text{ GeV}^2$ and the associated one-sigma uncertainty band for both (a) NLO and (b) NNLO distributions.

of effect as data set normalisations sometimes floating to one end of the allowed uncertainty. The combination $4/9(u + \bar{u}) + 1/9(d + \bar{d} + s + \bar{s})$ is very well constrained by structure function data. The fit to this can be obtained with the minimum momentum of partons if there is a large proportion of up and antiup quarks, so in the absence of higher order corrections the fit tends to choose this scenario if given the choice, and hence push the strange proportion down. Indeed it was recently noted [159] that if the strange distribution is left completely free in a fit to structure function data alone it settles on a marginally negative value (we have also observed this phenomenon), albeit with very large uncertainty. We remove this tendency from the fit by fixing the branching ratio at $B_c = 0.099$ due to complete correlation with the $(s + \bar{s})$ normalisation.

8.2 Impact on strange and antistrange distributions

In Fig. 35 we show the ratio of the strange sea to the non-strange sea,

$$\kappa(x, Q^2) \equiv \frac{s(x, Q^2) + \bar{s}(x, Q^2)}{\bar{u}(x, Q^2) + \bar{d}(x, Q^2)}, \quad (81)$$

at the input scale $Q_0^2 = 1 \text{ GeV}^2$ at both NLO and NNLO. At small x this suppression factor is just a little below the previous default values of 0.5 at NLO and 0.44 at NNLO. However, there is a strong suggestion of additional suppression of the strange sea at large x relative to the non-strange sea. This is as one might expect if the suppression is due to the mass of the strange quark since at high x one is nearer to the production threshold. As seen in Section 6.5, and illustrated in Fig. 20(a), the suppression of the strange distribution is consistent with the expectation from considering the strange quark to have a mass $m_s^2 \sim 0.1 \text{ GeV}^2$. In Fig. 36 we also show the difference between s and \bar{s} at the input scale. Note that s and \bar{s} are constrained by the dimuon data for $0.01 \lesssim x \lesssim 0.2$, the region where the central value of the asymmetry is non-zero. However, note that the asymmetry is non-zero just outside the one-sigma uncertainty

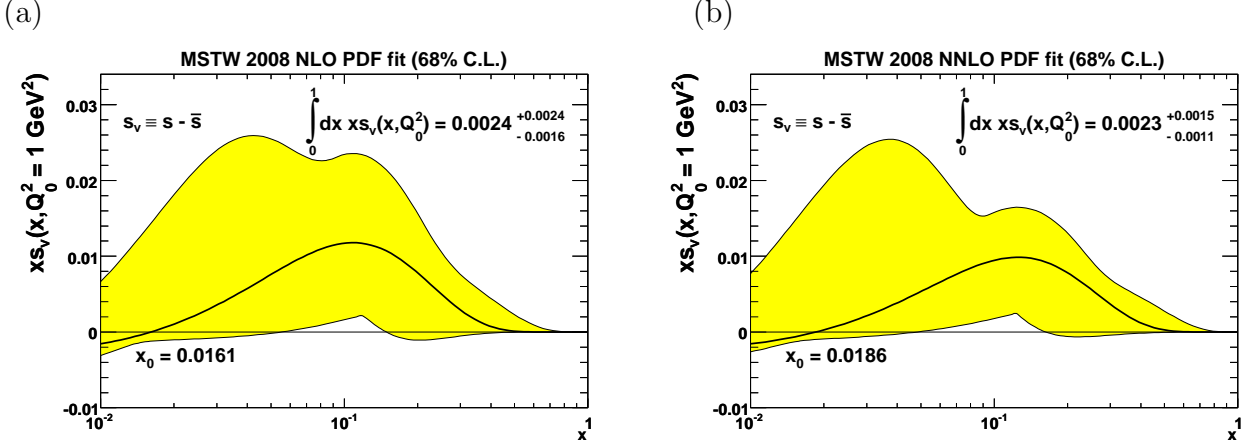


Figure 36: The strange asymmetry, $xs - x\bar{s}$, at the input scale $Q_0^2 = 1 \text{ GeV}^2$ and the associated one-sigma uncertainty band for both (a) NLO and (b) NNLO distributions.

Q^2/GeV^2	$\int_0^1 dx x(s - \bar{s})$			$\int_0^1 dx x(s + \bar{s}) / \int_0^1 dx x(\bar{u} + \bar{d})$		
	LO	NLO	NNLO	LO	NLO	NNLO
1	$0.0026^{+0.0022}_{-0.0017}$	$0.0024^{+0.0024}_{-0.0016}$	$0.0023^{+0.0015}_{-0.0011}$	$0.24^{+0.03}_{-0.03}$	$0.33^{+0.04}_{-0.04}$	$0.36^{+0.03}_{-0.04}$
10	$0.0019^{+0.0016}_{-0.0012}$	$0.0018^{+0.0018}_{-0.0012}$	$0.0016^{+0.0011}_{-0.0009}$	$0.45^{+0.02}_{-0.03}$	$0.53^{+0.02}_{-0.03}$	$0.54^{+0.02}_{-0.03}$
100	$0.0016^{+0.0014}_{-0.0010}$	$0.0015^{+0.0015}_{-0.0010}$	$0.0013^{+0.0010}_{-0.0007}$	$0.55^{+0.02}_{-0.02}$	$0.62^{+0.02}_{-0.02}$	$0.63^{+0.02}_{-0.02}$

Table 7: Integrated strangeness momentum asymmetry and integrated strangeness fraction for $Q^2 = 1, 10, 100 \text{ GeV}^2$. Note that $\int_0^1 dx x(s - \bar{s}) \approx 0.007$ (for typical $Q^2 \sim 20 \text{ GeV}^2$) is required to bring the NuTeV $\sin^2 \theta_W$ to the world average value. The one-sigma PDF uncertainties are given.

band, and it is certainly consistent with zero within the 90% C.L. uncertainty band. As with the total strange distribution the plots for the NLO and NNLO analyses are very similar. Indeed, this is quantified in Table 7 which shows the integrated strangeness momentum asymmetry,

$$\int_0^1 dx x [s(x, Q^2) - \bar{s}(x, Q^2)], \quad (82)$$

and the relative strength of the strange quark sea

$$\frac{\int_0^1 dx x [s(x, Q^2) + \bar{s}(x, Q^2)]}{\int_0^1 dx x [\bar{u}(x, Q^2) + \bar{d}(x, Q^2)]}, \quad (83)$$

for $Q^2 = 1, 10, 100 \text{ GeV}^2$, for the LO, NLO and NNLO parton sets. Notice that in both cases there is good perturbative convergence, though the LO strangeness fraction is a little low compared to the higher orders. This was not guaranteed to be the case, and relies on a careful definition of the GM-VFNS for charged-current structure functions, as outlined in Section 4.

The uncertainty on the s and \bar{s} distributions has increased significantly due to this new treatment, since previously it was simply the same (proportionally) as for the much more accurately determined \bar{u} and \bar{d} distributions. This increase is illustrated clearly in Fig. 20(b), and is one of the main reasons for the increase in uncertainty for the total sea seen in Fig. 19. The increased strange uncertainty also leads to a slight increase in the uncertainty for the \bar{u} and \bar{d} distributions. From constraints on the different weighting in the sum of sea quarks, mainly on the charge-weighted combination in neutral-current DIS, there is an anti-correlation between strange and \bar{u}, \bar{d} . However, this is a fairly small effect, and is difficult to disentangle from the consequences of other changes in our procedure (e.g. accounting for data set normalisation uncertainties). However, as well as an understanding of the uncertainties we now have much more confidence in our central values for the strange and antistrange distributions. The precise values are important for current experiments since the strange quark and antiquark distributions make significant contributions to W^+ and W^- production at hadron colliders, particularly at the LHC, through $c\bar{s}$ and $\bar{c}s$ partonic fusion. The W^\pm asymmetry is driven mainly by the difference between $u\bar{d}$ and $\bar{u}d$ fusion, but also has some sensitivity to the strange quark asymmetry. It is worth noting that the associated production processes $pp \rightarrow W^- + c + X$ and $pp \rightarrow W^+ + \bar{c} + X$ are a promising way to measure the strange quark distributions, where the LO partonic subprocesses are $gs \rightarrow W^-c$ or $g\bar{s} \rightarrow W^+\bar{c}$ [156]. Indeed, the Wc production process has recently been measured at the Tevatron [160].

8.3 Implications for the NuTeV $\sin^2 \theta_W$ anomaly

The determination of the momentum asymmetry $\int_0^1 dx x(s - \bar{s})$ of the s and \bar{s} distributions has important implications [161] for the anomaly in the measurement of $\sin^2 \theta_W$ reported by NuTeV [162]. The NuTeV extraction of $\sin^2 \theta_W$ from

$$R^- \equiv \frac{\sigma(\nu_\mu N \rightarrow \nu_\mu X) - \sigma(\bar{\nu}_\mu N \rightarrow \bar{\nu}_\mu X)}{\sigma(\nu_\mu N \rightarrow \mu^- X) - \sigma(\bar{\nu}_\mu N \rightarrow \mu^+ X)} \approx \frac{1}{2} - \sin^2 \theta_W \quad (84)$$

is about \sim three-sigma above the global average, and was at first thought to hint at new physics. However, there are two effects sensitive to parton distributions which should first be included: isospin violation ($u^p \neq d^n$, $d^p \neq u^n$) and the strange sea asymmetry ($s \neq \bar{s}$). Isospin violation gives

$$R^- \approx \frac{1}{2} - \sin^2 \theta_W + \left(1 - \frac{7}{3} \sin^2 \theta_W\right) \frac{\int_0^1 dx x [(u_v^p - d_v^n) - (d_v^p - u_v^n)]}{2 \int_0^1 dx x (u_v + d_v)}. \quad (85)$$

In fact, isospin violation is automatically generated by QED corrections to parton evolution, and was found by MRST [19] to remove slightly less than half of the total discrepancy in the NuTeV value of $\sin^2 \theta_W$ if evolution of the photon distribution (in the proton) from current quark mass scales was applied, while for evolution from constituent quark masses a reduction of about a quarter was obtained. The second effect is the possible existence of a momentum asymmetry $\int_0^1 dx x(s - \bar{s})$ of the strange sea. It gives

$$R^- \approx \frac{1}{2} - \sin^2 \theta_W - \left(1 - \frac{7}{3} \sin^2 \theta_W\right) \frac{\int_0^1 dx x(s - \bar{s})}{\int_0^1 dx x(u_v + d_v)}. \quad (86)$$

A positive value of $\int_0^1 dx \, x(s - \bar{s}) \sim 0.007$ [163, 164] would bring the NuTeV $\sin^2 \theta_W$ to the world average [132]. Indeed, we see that the analysis of the dimuon data gives a preference for a positive strange momentum asymmetry, which goes some way to further reduce the NuTeV $\sin^2 \theta_W$ anomaly. A reduction in the $\sin^2 \theta_W$ anomaly of about one-sigma is likely, but a reduction of up to two-sigma (or down to no effect) is within the range of the PDF uncertainties for the strange momentum asymmetry shown in Table 7.

9 Description of $F_2^{c\bar{c}}$ and $F_2^{b\bar{b}}$ data

9.1 Comparison to HERA data

The new $F_2^{c\bar{c}}$ data included in the MSTW fit [41–47] extend the data included in earlier MRST fits [41, 45]. In Fig. 37 we show the comparisons of our results at LO, NLO and NNLO to the charm structure function $F_2^{c\bar{c}}(x, Q^2)$ for all available H1 [41–44] and ZEUS [45–47] data. As one can see from the figure, and from the values of χ^2 in Table 2, the quality of the fit is good. At LO the fit undershoots badly for $Q^2 = 2 \text{ GeV}^2$, the absence of the $1/x$ term in the NLO coefficient function being most important in this region. However, the quick evolution at LO due to the large gluon and very large α_S quickly compensates and $F_2^{c\bar{c}}(x, Q^2)$ evolves most quickly at LO. The NNLO fit quality is best. This is clearly due to the improvement compared to NLO at the lowest x for $Q^2 = 2 \text{ GeV}^2$. At higher Q^2 the NNLO $F_2^{c\bar{c}}(x, Q^2)$ evolves most slowly, a feature that was highlighted as a consequence of heavy flavour at NNLO in Ref. [72]. This is most easily understood in the GM-VFNS, where the positive corrections to the coefficient functions at low Q^2 , and the negative contributions to the heavy flavour distributions from matching conditions influencing high Q^2 , both have the effect of flattening the slope. However, it has recently also been noticed in an approximate FFNS at NNLO (inclusion of threshold corrections and factorisation/renormalisation scale dependent terms) [95]. This slower evolution at NNLO is neither clearly preferred nor rejected by the data shown in Fig. 37, with some bins favouring NLO and some NNLO. Future more precise data may be more discriminating.

As outlined in Section 4, the heavy flavour treatment acquires some model dependence at NNLO, in the detailed form of the $\mathcal{O}(\alpha_S^3)$ coefficient functions, which are approximated using threshold and small- x limits. However, for charm production this term only appears with all scales fixed at $Q^2 = m_c^2$ in the fit to data (only depending on Q^2 for $Q^2 < m_c^2$), and is a small constant correction, which becomes increasingly less important at high Q^2 . The model dependence is mainly due to the treatment of the small- x limit. Reasonable variations of the parameters controlling this leads to variations in $F_2^{c\bar{c}}(x, Q^2)$ of approximately ± 0.03 for $x = 10^{-4}$ or ± 0.06 for $x = 10^{-5}$, and rather less than this for higher x , the variation becoming smaller than ± 0.01 at $x = 0.001$, and quickly dying away above this. This is enough to move the fit quality for $F_2^{c\bar{c}}(x, Q^2)$ data close to that at NLO in one direction, or to move the fit quality in the $Q^2 = 2 \text{ GeV}^2$ bin close to the best possible in the other direction. Because this is a correction to $F_2(x, Q^2)$ which is constant in Q^2 , this uncertainty has very little impact on the gluon distribution. The only real uncertainty it leads to is in the input for the sea distribution, which would change to compensate the Q^2 -constant change in $F_2(x, Q^2)$. The possible change is up to 4–5% for $x \lesssim 10^{-4}$, becoming much less at higher x . This is a lot smaller than the input uncertainty shown in Fig. 19(a). The relative uncertainty from this source would quickly become much smaller at high Q^2 , and would be a negligible influence on the uncertainty at $Q^2 = 100 \text{ GeV}^2$ shown in Fig. 19.

In Fig. 38 we show the predictions for the beauty structure function $F_2^{b\bar{b}}(x, Q^2)$ compared to the published H1 data [43, 44]. Clearly at all orders the comparison is good. It is also clear that these data currently provide no useful constraint on the gluon distribution. Since much of the data is in the vicinity of $Q^2 = m_b^2$, or not much higher, more precise data will be sensitive to the

Charm structure function, $F_2^{c\bar{c}}(x, Q^2)$

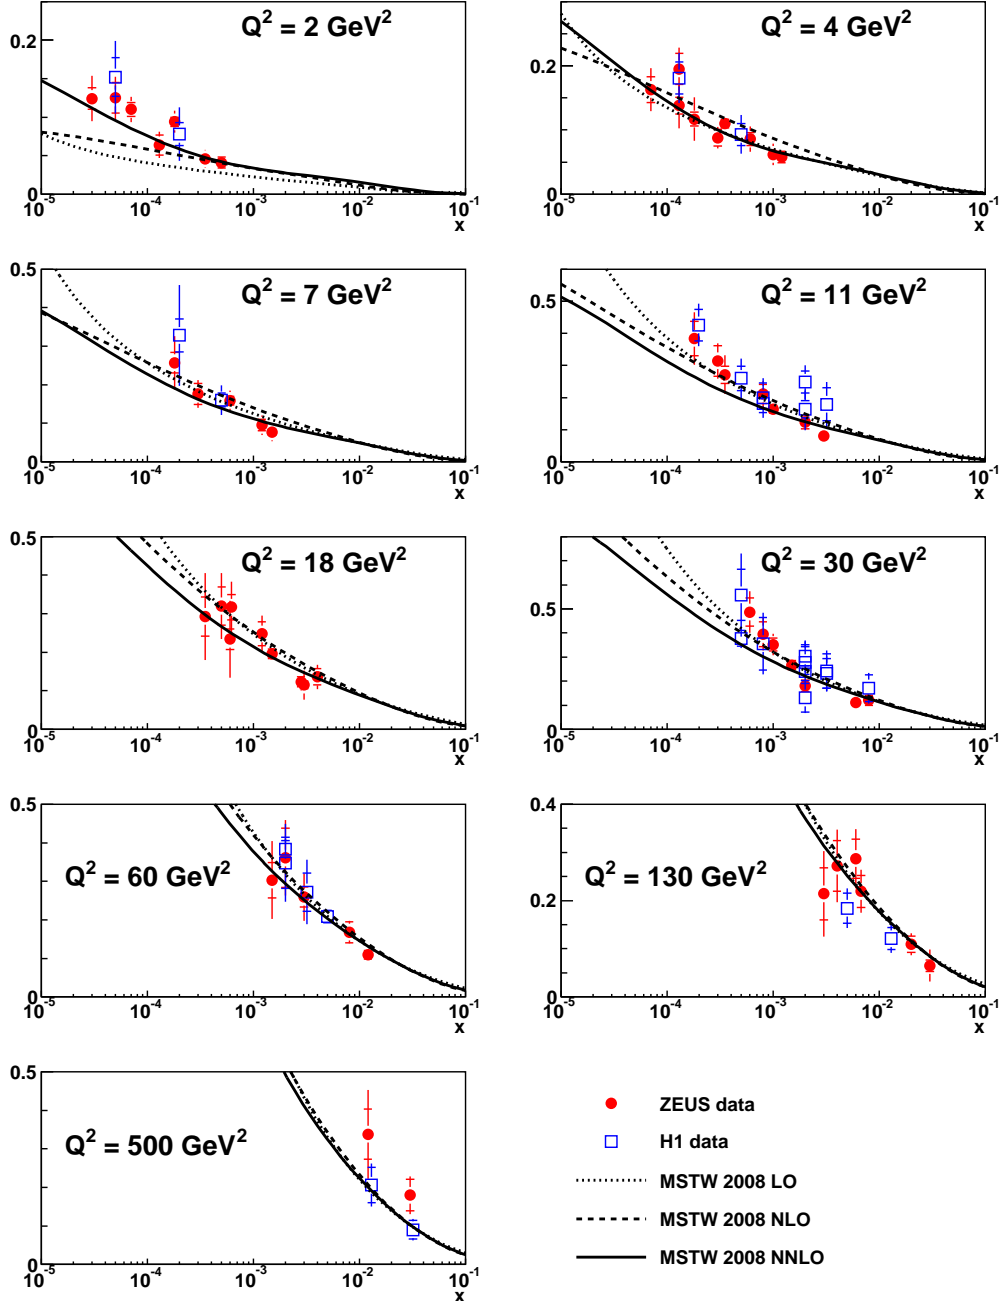


Figure 37: The charm structure function, $F_2^{c\bar{c}}(x, Q^2)$, compared to data from H1 [41–44] and ZEUS [45–47]. The Q^2 bins used correspond to the data points in Ref. [46]; the other data points have been shifted to these Q^2 values using the NNLO predictions.

Beauty structure function, $F_2^{b\bar{b}}(x, Q^2)$

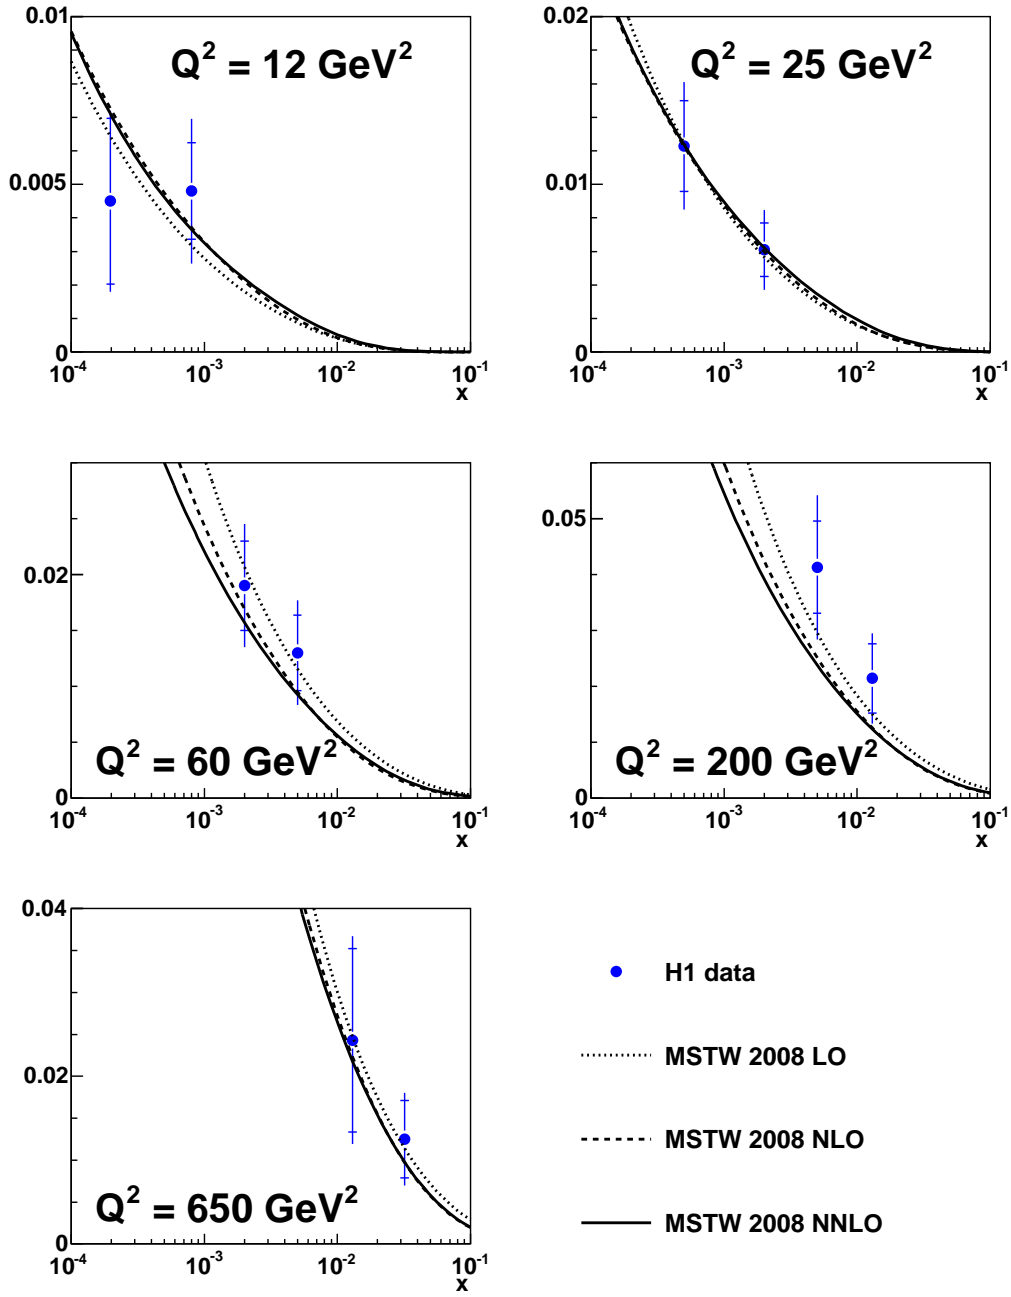


Figure 38: Predictions for the beauty structure function, $F_2^{b\bar{b}}(x, Q^2)$, compared to data from H1 [43, 44].

details of threshold contributions and in particular to the value of m_b rather than to the gluon.

9.2 Constraints on intrinsic charm

As explained in Section 4, our variable flavour number scheme allows for the possibility of including intrinsic charm with no theoretical problems. However, we simply choose not to include it in the default fit since we have no reason to believe it is anything other than a very small component of the total charm distribution, and none of the data in our global fit provide any constraint. The nonperturbatively generated contribution to the charm distribution is in general of $\mathcal{O}(\Lambda_{QCD}^2/m_c^2)$, but may be enhanced at large x [98], whereas all data on $F_2^{c\bar{c}}(x, Q^2)$ from HERA are at small x .

In this section we use the only data which currently do show any sensitivity, the EMC data on $F_2^{c\bar{c}}(x, Q^2)$ [165]. In Fig. 39 we show the comparison of the EMC charm data to our predictions at NLO and NNLO. In general the comparison to data is quite poor, tending to overshoot at the lower values of x and Q^2 . This is worse at NNLO than at NLO because the NNLO threshold corrections cause a significant increase. Hence, at fixed NLO and NNLO with $m_c = 1.4$ GeV there is some room for intrinsic charm at $x > 0.2$, but much more tendency to overshoot data at smaller x . This latter problem must be addressed before drawing conclusions about the intrinsic charm contribution. When considering this comparison it is pertinent to remember that when calculating heavy flavour structure functions we use $W^2 = 4m_c^2$ as the threshold, whereas in reality we need to produce mesons, i.e. W^2 is really required to be a little greater. In practice this is a higher-twist effect, but the EMC data are not so far from threshold, and are much more sensitive to these details than HERA data. Hence, we consider a comparison where we replace m_c^2 by $m_c^2(1 + \Lambda^2/m_c^2)$ in the threshold dependent parts of coefficient functions, where Λ^2 is a binding energy which we take to be $\Lambda = 0.2$ GeV. We make no change in the off-threshold parts of coefficient functions, nor in PDF definitions. The result is labelled as NLO MT in Fig. 39. It clearly matches lower Q^2 EMC data better, and the prediction is slightly smaller at higher x (even with this modification NNLO predictions would generally be too large). The curve labelled NLO MT + IC includes a 0.3% integrated number density contribution of intrinsic charm (and the same for anticharm) using the parameterisation in Ref. [98] with the upper limit on x modified to the correct threshold value. We note that this is only $\approx 1/10$ of the upper limit obtained by CTEQ [99]. This seems to be about the maximum that can be included even with nonperturbative suppression of standard fixed-order contributions, and even then only at NLO. A much larger contribution would overshoot most data at $x > 0.1$.²¹ Hence, if the EMC data are to be believed, there is no room for a very sizeable intrinsic charm contribution, although there is a suggestion of some deficit at the highest x values.

²¹Note that 0.3% was presented as the most likely value in 1983 using very old PDFs and only the LO perturbative contribution [166].

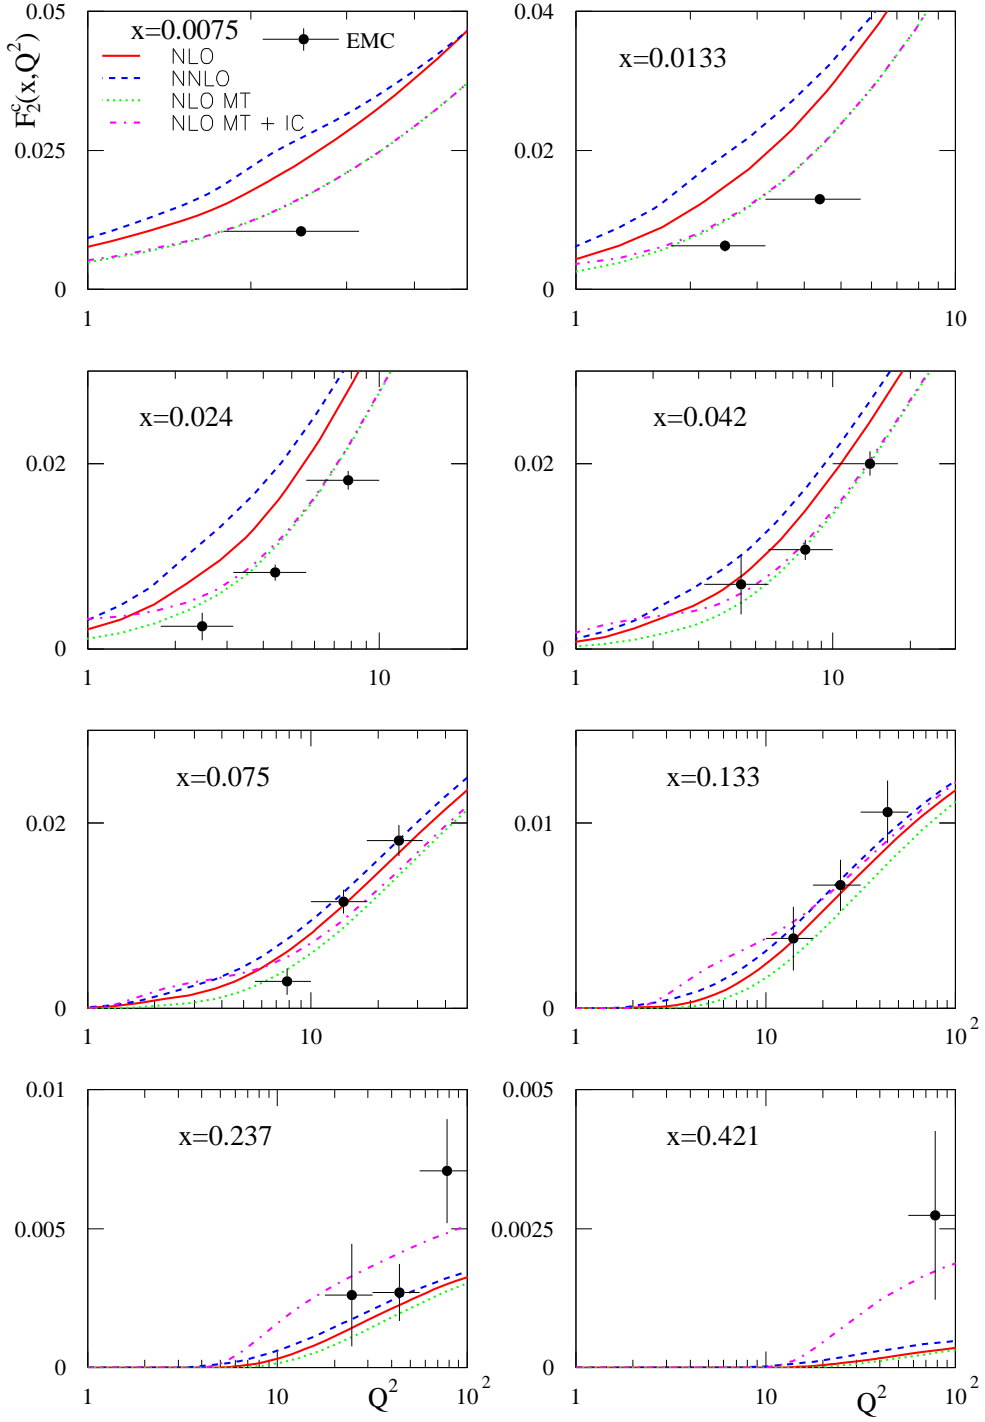


Figure 39: The comparison of the EMC charm data [165] to our predictions at NLO and NNLO. MT stands for the modified threshold approach and IC stands for inclusion of intrinsic charm.

10 Description of low-energy Drell–Yan data

As in recent analyses [18, 21] we fit to the E866/NuSea Drell–Yan dilepton production data in pp collisions [107], now corrected for electromagnetic radiative corrections; moreover we do not include the pd collision data as we observe significant systematic differences between our predictions and these results. The Drell–Yan process is known up to NNLO for the rapidity (y) distribution [167]. However, the data are binned in the Feynman- x variable, $x_F = x_1 - x_2$. Previous MRST fits used LO kinematics ($p_T = 0$) when transforming between the two variables. However, the mean p_T of each data bin is available, and hence we now use the exact kinematics:

$$y = \ln \left(H + \sqrt{H^2 + 1} \right), \quad \text{where } H \equiv \frac{x_F \sqrt{s}}{2\sqrt{M^2 + p_T^2}}. \quad (87)$$

At LO a good fit to the data is impossible while simultaneously fitting to structure function data which probe the same range of x and hard scale. This is because there is a large difference in the size of the corrections to the respective coefficient functions when going from the spacelike to the timelike regime, i.e. there is a factor of $1 + (\alpha_S(M)/\pi)C_F\pi^2/2$ for Drell–Yan production at NLO. Since this is missing at LO, the normalisations of the two calculated quantities cannot be made consistent with the structure function data and Drell–Yan data at the same time. Previous MRST fits used a constant factor of 1.3 to enhance the LO Drell–Yan cross section, but we have improved this and now use a K -factor consisting of the π^2 -enhanced term. At NLO and NNLO we isolate the α_S dependence from the K -factor by writing:

$$K^{\text{LO}}(M) = 1 + \left(\frac{\alpha_S(M)}{\pi} \right) \frac{C_F\pi^2}{2}, \quad (88)$$

$$K^{\text{NLO}}(M, y) = 1 + \left(\frac{\alpha_S(M)}{\pi} \right) C(M, y), \quad (89)$$

$$K^{\text{NNLO}}(M, y) = 1 + \left(\frac{\alpha_S(M)}{\pi} \right) D(M, y) + \left(\frac{\alpha_S(M)}{\pi} \right)^2 E(M, y). \quad (90)$$

We calculate the C -, D - and E -factors with the DYRAP program [167] using NLO or NNLO PDFs at the respective order. This method, first used in Ref. [21], is a considerable improvement on previous MRST fits where only the overall $K(M, y)$ was calculated for a fixed value of α_S . (Additionally, NNLO fits prior to Ref. [21] used only NLO K -factors.)

The quality of the fit to the E866/NuSea data is illustrated in Fig. 40 at NNLO. As seen in Table 2, the fit quality is very similar at LO and NLO and the figures of data against theory would be similar to that at NNLO. As one can see, there is no systematic failure in the comparison between data and theory — the relatively large χ^2 value is due to a large scatter in the data points. The only exception is that in the lowest mass bin the data lie somewhat above the theory at low rapidities, although low M is where additional theoretical corrections are most likely to play some role and also where there may be complications in the mass reconstruction from the proximity to the J/ψ and Υ resonances. Additionally, unlike alternative studies [168], we do not see any significant inconsistency between the E866/NuSea pp data and the rest of the global

Ratio of E866/NuSea pp Drell-Yan data ($\times 1.087$) to
MSTW 2008 NNLO PDF fit, $\chi^2 = 237$ for 184 pts.

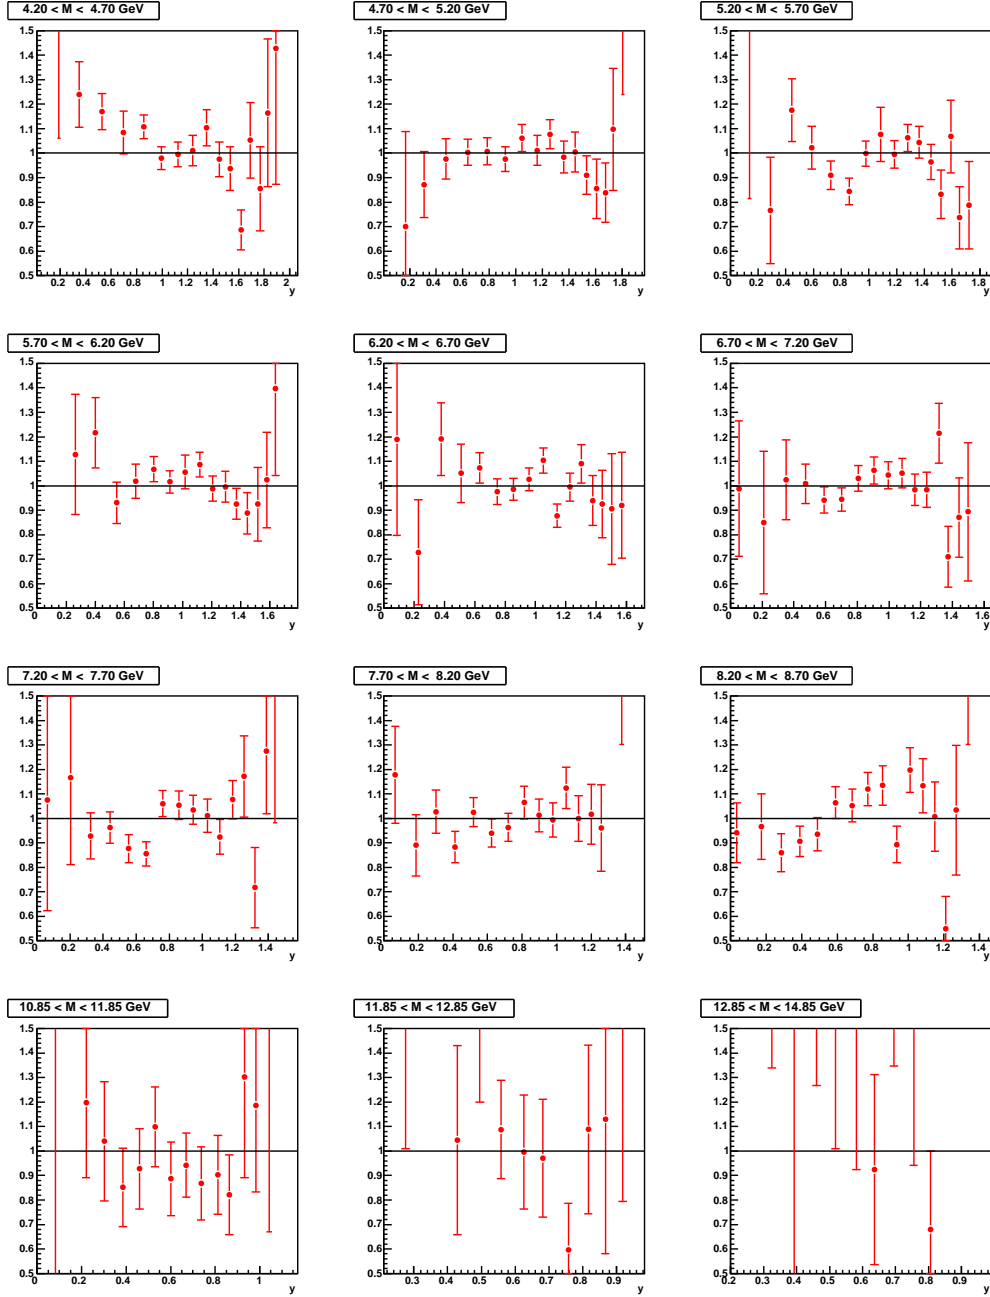


Figure 40: The ratio of E866/NuSea pp Drell-Yan data points to the predictions of the NNLO PDF fit as a function of rapidity for different mass bins.

fit. Our χ^2 for the remaining data only decreases by a few units when these Drell–Yan data are removed, i.e. by much less than the tolerance on any of our eigenvectors. We do note, however, that the positive corrections of order 10% to the cross section at NNLO require the data to move up to a normalisation of 1.087, i.e. slightly outside the quoted uncertainty of 6.5%.

As in previous analyses we include the E866/NuSea pd/pp ratio data [108], which constrains the $\bar{d} - \bar{u}$ difference. The experimental extraction of the pd/pp ratio by E866/NuSea used a larger data set than the individual cross section measurements, although there was some overlap. Moreover, the analysis was of a different form using different cuts and averages over much larger data subsets, and therefore had little sensitivity to those lower statistics regions where the discrepancies with the pd data occur, in particular at the highest x_F values. We have improved the theoretical treatment of the pd/pp ratio compared to previous analyses by calculating the separate K -factors for pp and pd collisions. This has an effect of at most 2%, and does aid the quality of the global fit very slightly. There is no significant change in our results for the difference $\bar{d} - \bar{u}$, though the uncertainty is reduced a little compared to previous analyses, as seen in Fig. 18.

11 W and Z boson production at the Tevatron

W and Z boson production at high-energy hadron colliders has traditionally provided important constraints on parton distributions. While low-energy Drell–Yan production directly probes the quark sea, W and Z production (via $q\bar{q} \rightarrow V$, with $V = W^\pm, Z$) at proton–antiproton colliders are mainly sensitive to the valence u and d quarks, in different combinations to structure function measurements in deep-inelastic scattering.

Because of the overall experimental luminosity uncertainty (of order $\pm 6\%$ at the Tevatron), the *total* W and Z cross section measurements do not provide competitive constraints on the PDFs, although they do provide an important cross-check on the parton distributions and on the hard-scattering framework (see Section 15 below). The shape of the W and Z rapidity distributions, however, is sensitive to the shape of the parton distributions, and *does* provide a powerful constraint. We have for many years used CDF data on the $W \rightarrow \ell\nu$ charge asymmetry [51] in our global fits, and in the following subsection we study the impact of recent CDF and DØ measurements in the updated MSTW fit. There are also nowadays very precise Tevatron data on the Z rapidity distribution, which provides complementary information to the $W \rightarrow \ell\nu$ charge asymmetry measurement. We include these data in the global fit for the first time, and study their consequences.

As for the Drell–Yan process, the W and Z rapidity distributions have been calculated to NNLO [167], and so we are able to perform consistent fits at LO, NLO and NNLO.

11.1 Description of Tevatron W asymmetry data

The W charge asymmetry at the Tevatron is defined by

$$A_W(y_W) = \frac{d\sigma(W^+)/dy_W - d\sigma(W^-)/dy_W}{d\sigma(W^+)/dy_W + d\sigma(W^-)/dy_W} \approx \frac{u(x_1)d(x_2) - d(x_1)u(x_2)}{u(x_1)d(x_2) + d(x_1)u(x_2)}, \quad (91)$$

where $x_{1,2} = (M_W/\sqrt{s}) \exp(\pm y_W)$. This description in terms of the PDFs is valid at leading order when sea-quark contributions are neglected. In practice, it is usually the lepton charge asymmetry which is measured, defined in a similar way as

$$A(\eta_\ell) = \frac{d\sigma(\ell^+)/d\eta_\ell - d\sigma(\ell^-)/d\eta_\ell}{d\sigma(\ell^+)/d\eta_\ell + d\sigma(\ell^-)/d\eta_\ell}, \quad (92)$$

where η_ℓ is the pseudorapidity of the charged lepton. Defining the emission angle of the charged lepton relative to the proton beam in the W rest frame by $\cos^2 \theta^* = 1 - 4E_T^2/M_W^2$ leads to

$$y_\ell = y_W + \frac{1}{2} \ln \left(\frac{1 + \cos \theta^*}{1 - \cos \theta^*} \right). \quad (93)$$

Care must be taken when using a valence-quark-only approximation to the lepton asymmetry, since near the edges of phase space, i.e. $\cos \theta^* \approx \pm 1$, sea-quark contributions can become

important. In fact, neglecting overall factors, the numerator of (92) can be approximated by²²

$$u(x_1)d(x_2)(1 - \cos \theta^*)^2 + \bar{d}(x_1)\bar{u}(x_2)(1 + \cos \theta^*)^2 - d(x_1)u(x_2)(1 + \cos \theta^*)^2 - \bar{u}(x_1)\bar{d}(x_2)(1 - \cos \theta^*)^2, \quad (94)$$

where for $\cos \theta^* \approx 1$ the leading $(\bar{d}\bar{u})$ sea–sea contribution for W^+ production is enhanced relative to the valence–valence contribution by the large $(1 + \cos \theta^*)^2$ term arising from the $V - A$ decay to leptons. Hence if the lepton charge asymmetry is measured in different bins of lepton E_T , the sensitivity to antiquark PDFs is greatest at lower E_T . For example, if $E_T = 25$ GeV then $\cos \theta^* = \pm 0.78$ and, taking the positive value, then from (94) the antiquark contribution will be enhanced relative to the quark contribution by a factor $(1 + \cos \theta^*)^2 / (1 - \cos \theta^*)^2 = 68$. In this case, in the forward direction, leptons from W^+ produced from quarks will be mainly from bosons with rapidity $y_W = y_\ell + 1.05$ (taking $\cos \theta^* = -0.78$), while leptons from W^+ produced from antiquarks will be mainly from bosons with rapidity $y_W = y_\ell - 1.05$ (taking $\cos \theta^* = +0.78$). At high $\eta_\ell \simeq y_\ell$ (not too close to the kinematic limit) this can lead to the latter being a significant proportion, or even dominating. Hence, in addition to measuring the overall shapes of the quark distributions, the lepton charge asymmetry also probes the separation into valence and sea quarks. In practice, this is particularly so for the less well-constrained down quark.

In the global fit, we of course include all quark and gluon contributions to the appropriate order in perturbation theory. We calculate K -factors for the differential cross sections $d\sigma(\ell^\pm)/d\eta_\ell$ using the FEWZ code [73] with LO, NLO or NNLO PDFs at the respective order.²³ We also include the effect of a finite W width in these K -factors.²⁴

In Fig. 41 we show the effect of a finite W width, NLO corrections and NNLL p_T^W -resummation on the lepton pseudorapidity distribution when using a common set of parton distributions (MRST 2004 NLO). Then in Fig. 42 we show the corresponding predictions for the lepton charge asymmetry. In the lower E_T^e bin, finite W width and higher-order effects are important only in the forward direction, meaning that these features have a significant effect on the asymmetry shown in Fig. 42. (Recall that CP invariance requires that $d\sigma(\ell^+)/d\eta_\ell = d\sigma(\ell^-)/d(-\eta_\ell)$.) In the upper E_T^e bin, these features are almost symmetric about $\eta_e = 0$, meaning that the effect on the asymmetry is small. From Fig. 42 we see that at NLO the effect on the asymmetry of the finite W width is smaller than the experimental uncertainties, but is not insignificant, and hence we include this. The further effect of resummations is generally smaller still, so given that our fit is meant to correspond to PDFs at a fixed perturbative order we do not include the effect of resummations in our K -factors.

²²Contributions involving s , c and b quarks are not displayed, as these are very small at Tevatron energies.

²³We find that the VEGAS Monte Carlo integration in FEWZ does not converge at NNLO for lepton production for bin widths of $\Delta\eta_\ell = 0.2$ to sufficient accuracy of $\mathcal{O}(1\%)$ within a reasonable run time of $\mathcal{O}(\text{months})$. This problem has been noticed by other groups (see, for example, Refs. [169, 170]). Since the effect of NNLO corrections is expected to cancel almost entirely when calculating the asymmetry $A(\eta_\ell)$, we instead use NLO K -factors calculated with NNLO PDFs. We have confirmed that the NNLO corrections to the W asymmetry (as a function of the W rapidity) are extremely small; see Fig. 11 of Ref. [167].

²⁴The future use of APPLGRID [171] may eliminate the need to resort to calculating LO W and Z cross sections multiplied by fixed K -factors during PDF fits, although the latter method should still be a very good approximation.

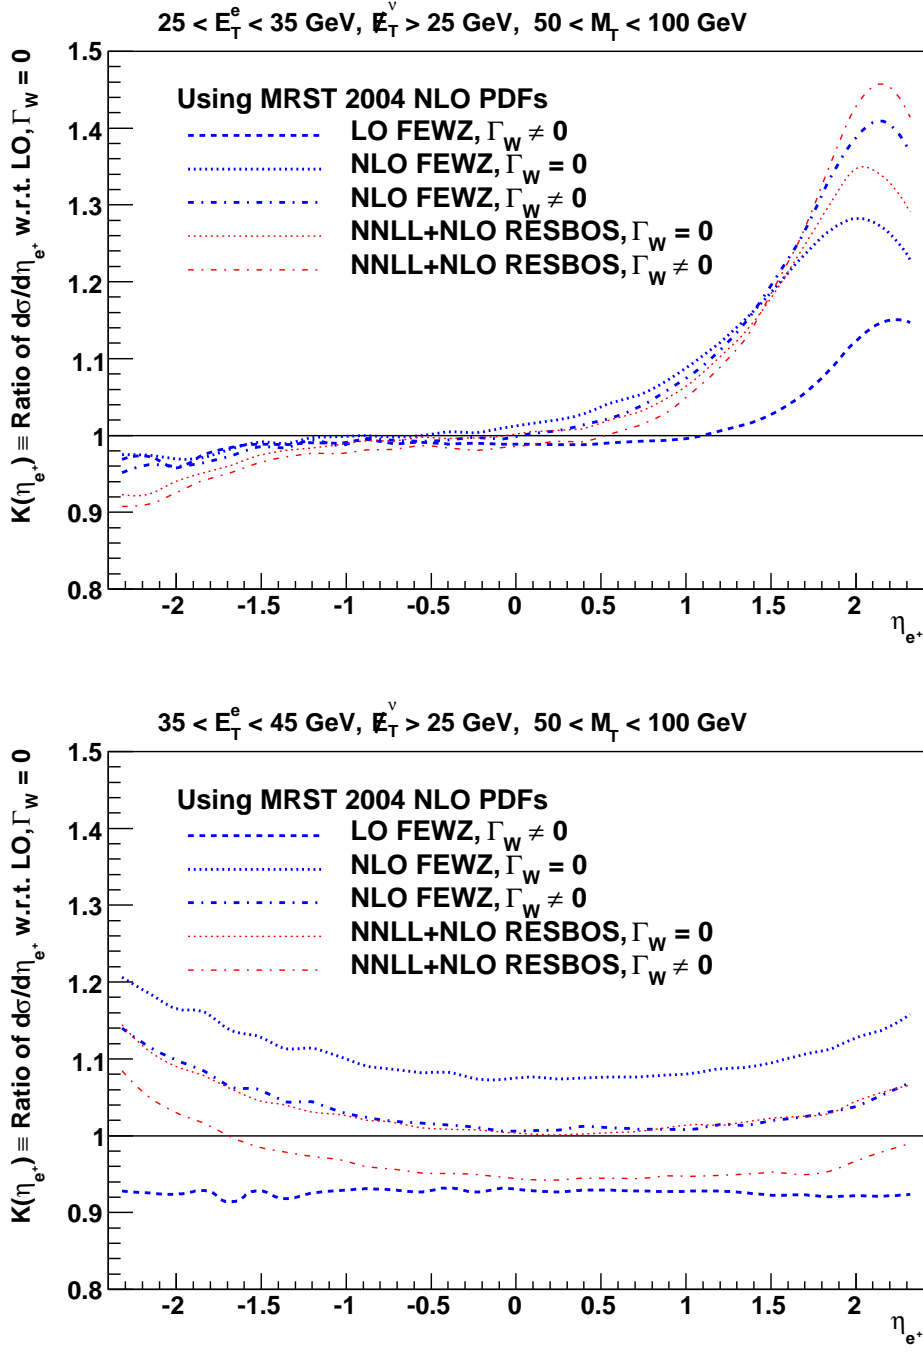


Figure 41: $d\sigma/d\eta_{e^+}$ K -factors for CDF kinematic cuts [48] for various methods of calculation.




Figure 42: $W \rightarrow e\nu$ charge asymmetry from CDF [48] for various methods of calculation.

CDF data on lepton charge asymmetry from $W \rightarrow e\nu$ decays

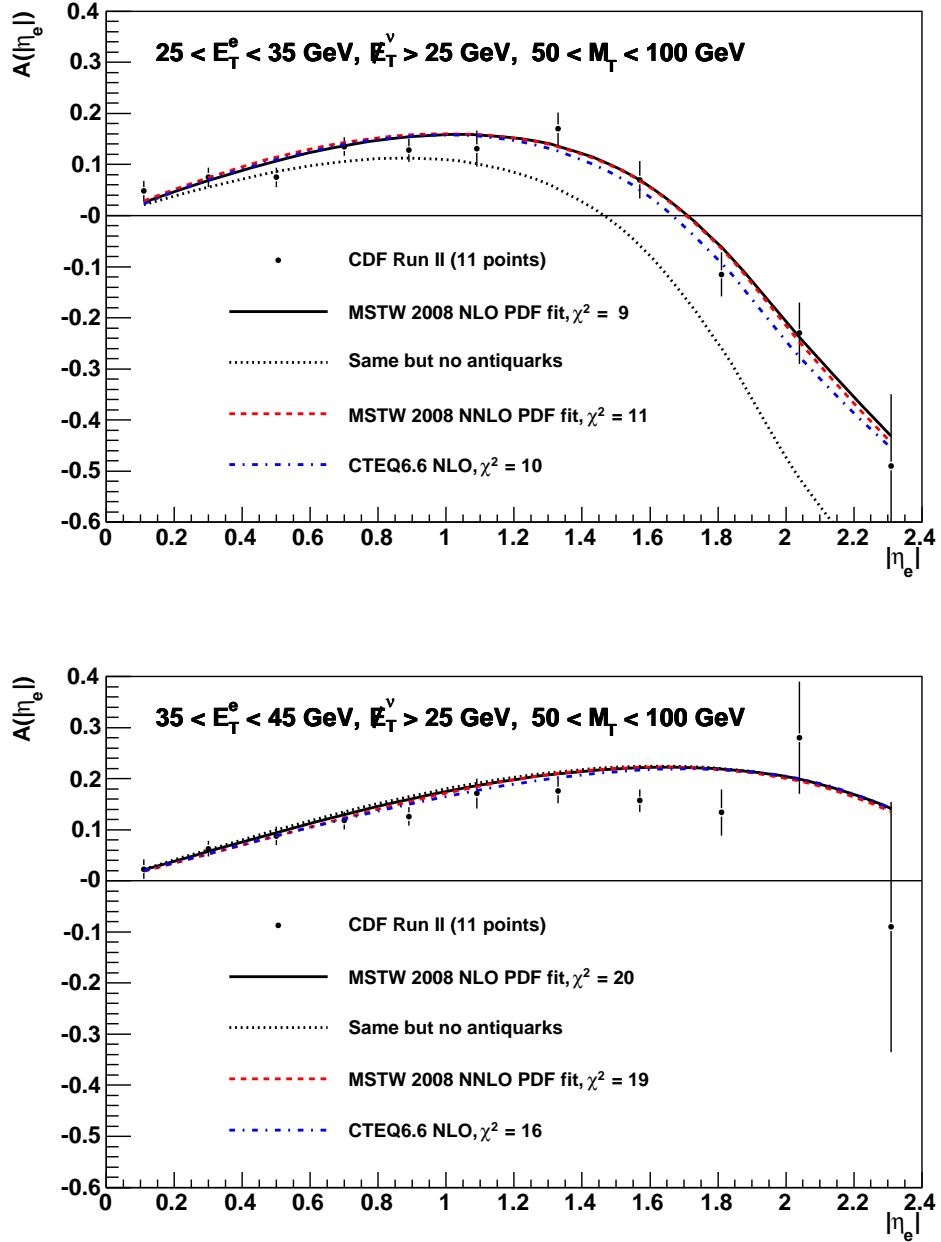


Figure 43: A comparison of theory and data for $W \rightarrow e\nu$ charge asymmetry from CDF [48].




Figure 44: A comparison of theory and data for $W \rightarrow \mu\nu$ charge asymmetry from DØ [49].

Figure 43 shows the fit to the most recent CDF $W \rightarrow e\nu$ lepton charge asymmetry data [48]. The data are presented in two bins of lepton E_T . In each case the data are reasonably well fitted at both NLO and NNLO. Also shown is the NLO fit with the antiquark contributions removed. The difference is much more significant in the lower E_T bin, as expected. In the higher E_T bin there appears to be some evidence for a systematic discrepancy between the fit and the data, particularly at high $|\eta_e|$ (these data seem to be described slightly better by CTEQ6.6 PDFs). This could in principle be removed by increasing the d quark distribution in the $x \sim 0.1$ region, thereby enhancing the negative $d(x_1)$ contribution in the numerator in (94), but this would produce a tension with DIS structure function data. The overall impact on the d -quark distribution will be discussed further below.

The corresponding fits to the DØ $W \rightarrow \mu\nu$ lepton charge asymmetry data from Ref. [49] are shown in Fig. 44. The data are presented in a single lepton $p_T^\mu > 20$ GeV bin, and the antiquark contributions are therefore sizeable. Some systematic discrepancy at high $|\eta_\mu|$ is again apparent. Note that there are more recent DØ electron data [50] than the muon data shown in Fig. 44. However when we include these DØ electron data in the global analysis we are unable to obtain a good quality fit, with significant tension between the new DØ electron asymmetry data and the DIS structure function data (both proton and deuterium) and with low-mass Drell–Yan data. Furthermore, we notice considerable tension between the DØ electron asymmetry data [50] and the preliminary CDF W asymmetry data [172].²⁵ Pending further investigation of these issues, we have decided not to include either the more recent DØ data [50] or the CDF W asymmetry

²⁵The CDF W asymmetry data have since been published [173].

data [172, 173] in the current fit.

11.2 Description of Tevatron Z rapidity distribution data

The integrated luminosity for Run II at the Tevatron has allowed a high-precision measurement of the Z rapidity distribution. Analogously to the low-energy Drell–Yan distribution $d^2\sigma/dMdy$ discussed in Section 10, this provides a constraint on the (mainly quark) PDFs over a broad range of x at $Q^2 \sim M_Z^2$. The mix of quark distributions probed is approximately

$$0.29 u(x_1) u(x_2) + 0.37 d(x_1) d(x_2), \quad (95)$$

in contrast to the $(4u\bar{u} + d\bar{d})$ combination probed in pp Drell–Yan production. The Z rapidity distribution is therefore particularly sensitive to the d distribution, providing complementary information to the W charge asymmetry discussed above.

The theoretical calculations are performed using the same form of K -factors as for the low-mass Drell–Yan data (88), (89) and (90), with the NLO and NNLO coefficients calculated exactly using the VRAP program [167].

In Fig. 45 we show the results of the NLO and NNLO fits compared to the Z/γ^* rapidity shape distribution measured in the mass range $71 < M_{ee} < 111$ GeV by DØ using 0.4 fb^{-1} of Run II data [53]. In Fig. 46 we show the results of the NNLO fit compared to the preliminary Z/γ^* rapidity distribution measured in the mass range $66 < M_{ee} < 116$ GeV by CDF using 2.1 fb^{-1} of Run II data [52]. The data are well fitted, particularly when the systematic uncertainties are taken into account. There is evidence of a slight excess of data over theory at high rapidity ($y > 2.2 \Rightarrow x_1 > 0.4$ and $x_2 < 0.005$).

During the fit the Z rapidity distribution was calculated at LO in the narrow-width approximation ($M_{ee} = M_Z$) neglecting the γ^* contribution, then applying the appropriate K -factors to account for higher-order QCD corrections. To be more precise, the invariant mass range $66 < M_{ee} < 116$ GeV should be explicitly integrated over in the theory prediction allowing for a finite Z width and including the γ^* contribution. This more precise treatment has the effect of lowering the theoretical prediction by less than 1% over most of the rapidity range, rising above unity at $|y| \approx 2$, then rising more rapidly as the kinematic limit is approached at still larger rapidities. Refitting the PDFs with this additional factor included gives an increase in the fitted CDF normalisation by 1% and a better description of the CDF Z/γ^* rapidity data at large rapidity, with smaller shifts needed in the correlated systematic uncertainties. For the NLO fit, the χ^2 for the 29 CDF data points improves from 49 to 42, while the χ^2 for the 28 DØ data points improves from 19 to 17. However, the refitted PDFs are almost unchanged, with the differences very small compared to the quoted uncertainties. Since the CDF data are still unpublished and may change before publication, we will update our fits in the near future using the final Z/γ^* data and accounting for the precise M_{ee} range and γ^* contribution in the theoretical calculation.

Z/γ^* rapidity shape distribution from DØ

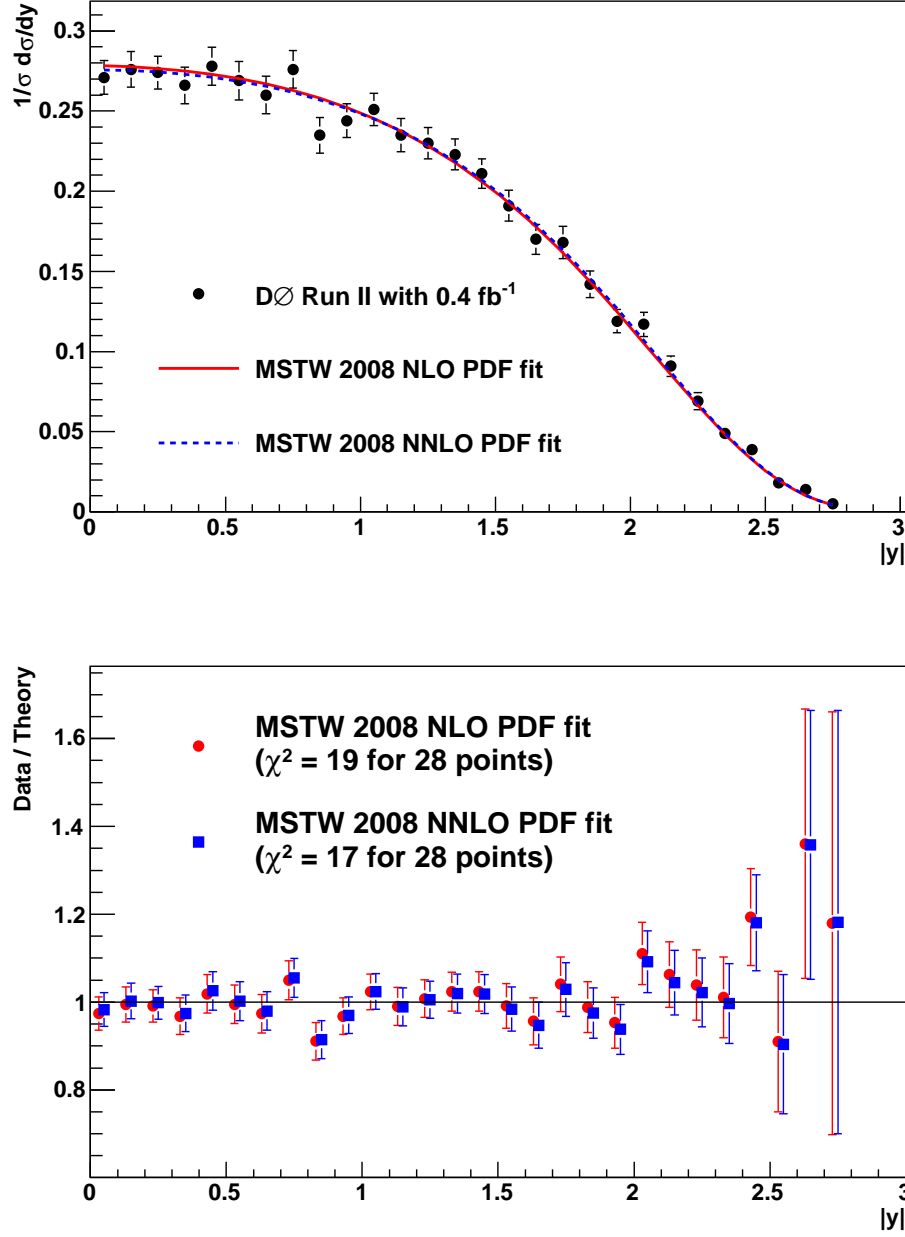


Figure 45: The comparison of data and theory for the Z/γ^* rapidity shape distribution from DØ [53].

Z/γ^* rapidity distribution from CDF

MSTW 2008 NNLO PDF fit, $\chi^2 = 50$ for 29 points





Figure 46: The comparison of data and theory for the Z/γ^* rapidity distribution from CDF [52].




Figure 47: The impact of new data on the down valence quark distribution at NNLO.

11.3 Impact on the down valence quark distribution

The W and Z data at the Tevatron are mainly sensitive to the up and down quark distributions of the proton. However, the up quark distribution is already well constrained by proton structure function data where it appears charge-weighted. Therefore, in practice, the Tevatron W and Z data mainly constrain the down quark distribution. In Fig. 47 we compare the down valence distribution at NNLO with that from the MRST 2006 NNLO set. Note that we now make a better choice of parameters in the input d_v distribution (7) to allow for more flexibility for the eigenvector PDF sets, which generally gives larger uncertainties than previously, despite the presence of more constraining data. Clearly there is a significant change in the shape, though the new and previous versions generally agree within the 90% C.L. limits. This change is mainly due to the inclusion of new Tevatron data, though there is also some influence from the new neutrino structure function data. The most significant change is the increase in $d_v(x, Q^2)$ for $x \sim 0.3$, seen more clearly in Fig. 61(b), which is influenced by the lepton asymmetry data. From the number sum rule this has to be compensated for by a decrease elsewhere, which is mainly for $x \sim 0.07$.

12 Description of jet data

Knowledge of the high- x gluon distribution is important for calculations of both signal and background in new physics searches at the LHC. The main direct constraint comes from Tevatron data on inclusive jet production. This is well illustrated by looking at the analyses of the Run I Tevatron jet data.

12.1 Description of Run I Tevatron jet data

Historically, early CDF Run I inclusive jet data at central rapidity [174] indicated an excess in the high- E_T jet spectrum using NLO QCD with the available parton distributions of the time. This discrepancy was initially taken as a possible indication of quark compositeness. It was found to be “impossible” [9] to describe simultaneously both the CDF jet data and the DIS structure function data with a common set of parton distributions having a physically motivated input form. However, it was found that the data could be accommodated in a global fit with an input gluon parameterisation with a somewhat contrived shoulder at high x [175]. Subsequently, DØ Run I jet data became available, which measured the E_T distribution at a range of rapidities [176]. NLO global analyses were able to fit all the jet data with physically reasonable forms, though there was some tension between the description of the high- E_T CDF and DØ data. This “saga” indicated the need to provide PDFs with associated uncertainties so that it can easily be seen in which regions of x the PDFs are relatively constrained or unconstrained. (These issues are further discussed in, for example, Refs. [177, 178].)

The Tevatron Run I jet data have been routinely included in the default global analyses since CTEQ4 [175] and MRST 2001 [14]. The MRST 2004 analysis [18] showed that a physical parameterisation of the input gluon distribution could be obtained by taking the input form (10) in the DIS scheme then transforming to the conventional $\overline{\text{MS}}$ scheme. The resultant $\overline{\text{MS}}$ parameterisation was found to automatically give a shoulder-like form at large x , which produced a better description of the Tevatron Run I inclusive jet data [176, 179] than the previous MRST 2001 PDFs.

At NLO, the cross section for inclusive jet production at the Tevatron can be written as

$$\sigma_{p\bar{p}} = \alpha_S^2(\mu_R^2) \sum_{a,b=q,g} [\hat{\sigma}_{ab}^{\text{LO}} + \alpha_S(\mu_R^2) \hat{\sigma}_{ab}^{\text{NLO}}] \otimes f_{a/p}(x_a, \mu_F^2) \otimes f_{b/\bar{p}}(x_b, \mu_F^2). \quad (96)$$

Usually, multidimensional (Monte Carlo) phase space integration is used to evaluate (96) with cuts on, for example, the transverse energy E_T and pseudorapidity η of the produced jets. However, this can take hours or days of CPU time, so it is impractical to include jet data in a PDF fit in this way. The old (approximate) solution used by MRST was to calculate “ K -factors”, defined as $\sigma_{p\bar{p}}^{\text{NLO}}/\sigma_{p\bar{p}}^{\text{LO}}$, for a given set of PDFs. Then six pseudogluon “data” points at $\mu_F^2 = 2000 \text{ GeV}^2$ (and a value of Λ_{QCD} corresponding to $\alpha_S(M_Z^2) \approx 0.115$) were inferred that would correspond to the best $\sigma_{p\bar{p}}$ needed to describe the data. These six pseudogluon points (with a free normalisation) and the Λ_{QCD} point were included in the MRST fit rather than the original jet data points. The comparison to the actual jet data points, including the inclusion of the correlated systematic uncertainties, was only made *after* the fit, though it was checked

that the χ^2 changes induced in the pseudodata were quantitatively similar to those in the final comparison to the jet data. The CTEQ group *do* include the jet data points explicitly in their fit, but the jet cross section is still calculated at LO and a NLO K -factor applied.

A much improved solution, which we adopt here, is to interpolate the PDFs and α_S so that they can be factorised from $\hat{\sigma}_{ab}$ in (96). Then the convolution integrals of (96) are replaced with multiplication. The phase space integration then only needs to be done once for a given set of kinematic cuts, with the result stored in a grid for later use during a PDF fit. This technique has previously been used by H1 [181] to determine the gluon density from DIS jet data, and by ZEUS [140] in a PDF fit including jet data from DIS and photoproduction. Grids corresponding to the kinematic cuts used in published Tevatron and HERA jet data are provided by the FASTNLO project [75] which uses the NLOJET++ calculations [76, 77]. (Another project with similar aims, which has also been applied to calculate W and Z production using MCFM [180], is the APPLGRID project [171].)

12.2 Description of Run II Tevatron jet data

Recently, Tevatron Run II inclusive jet data from both CDF [54, 55] and DØ [56] have become available with much higher statistics and much better control over the corrections necessary to go from the “raw” data to measurements at the parton “jet” level. For example, DØ now claim a relative uncertainty on the jet p_T of less than 2%. Our present global analysis uses all the available Run II jet data, including the complete information on correlated systematic uncertainties in the χ^2 definition (38). Due to the increased reliability, and statistics, of the Run II jet data, we now omit the Run I data from the global fit.

Another change is that we now use a renormalisation and factorisation scale $\mu_R = \mu_F = p_T$, rather than $\mu_R = \mu_F = p_T/2$.²⁶ The latter scale appears to be very convergent for central rapidities, but becomes unstable for larger y . For example, the NLO contribution is some 60–70% of that at LO for $y \sim 2.5$. In comparison, the scale choice $\mu_R = \mu_F = p_T$ leads to NLO corrections of ≈ 20 –30% for all y (except at the highest y and p_T , where they are smaller) and hence, while not as convergent near $y = 0$, the scale choice $\mu_R = \mu_F = p_T$ is reasonably stable for all rapidities of the data.²⁷

In Fig. 48 we show the comparison of the NLO fit with CDF Run II inclusive jet data using the k_T jet algorithm [54]. There is also an analysis of the same data using the cone-based Midpoint algorithm [55], but the k_T algorithm data [54] were published and implemented in our analysis first, therefore only the k_T algorithm data are included in the final global fit. Furthermore, there is a theoretical preference for using the k_T jet algorithm rather than a seeded cone algorithm, since the latter is not formally infrared safe, which in practice means that the perturbative calculation is only accurate to a given order of the expansion (NLO for the Midpoint algorithm). In Fig. 49 we show the comparison of the NLO fit to DØ Run II inclusive jet data using a cone jet algorithm [56]. We achieve an excellent description of both data sets shown in Figs. 48 and

²⁶Note that the Run II data are binned in terms of transverse momentum p_T and rapidity y , while the Run I data were binned in terms of transverse energy E_T and pseudorapidity η , although $E_T = p_T$ and $\eta = y$ for the massless jets used in Run I, but not for the massive jets used in Run II.

²⁷We thank M. Wobisch for bringing this to our attention.

CDF Run II inclusive jet data, $\chi^2 = 56$ for 76 pts.

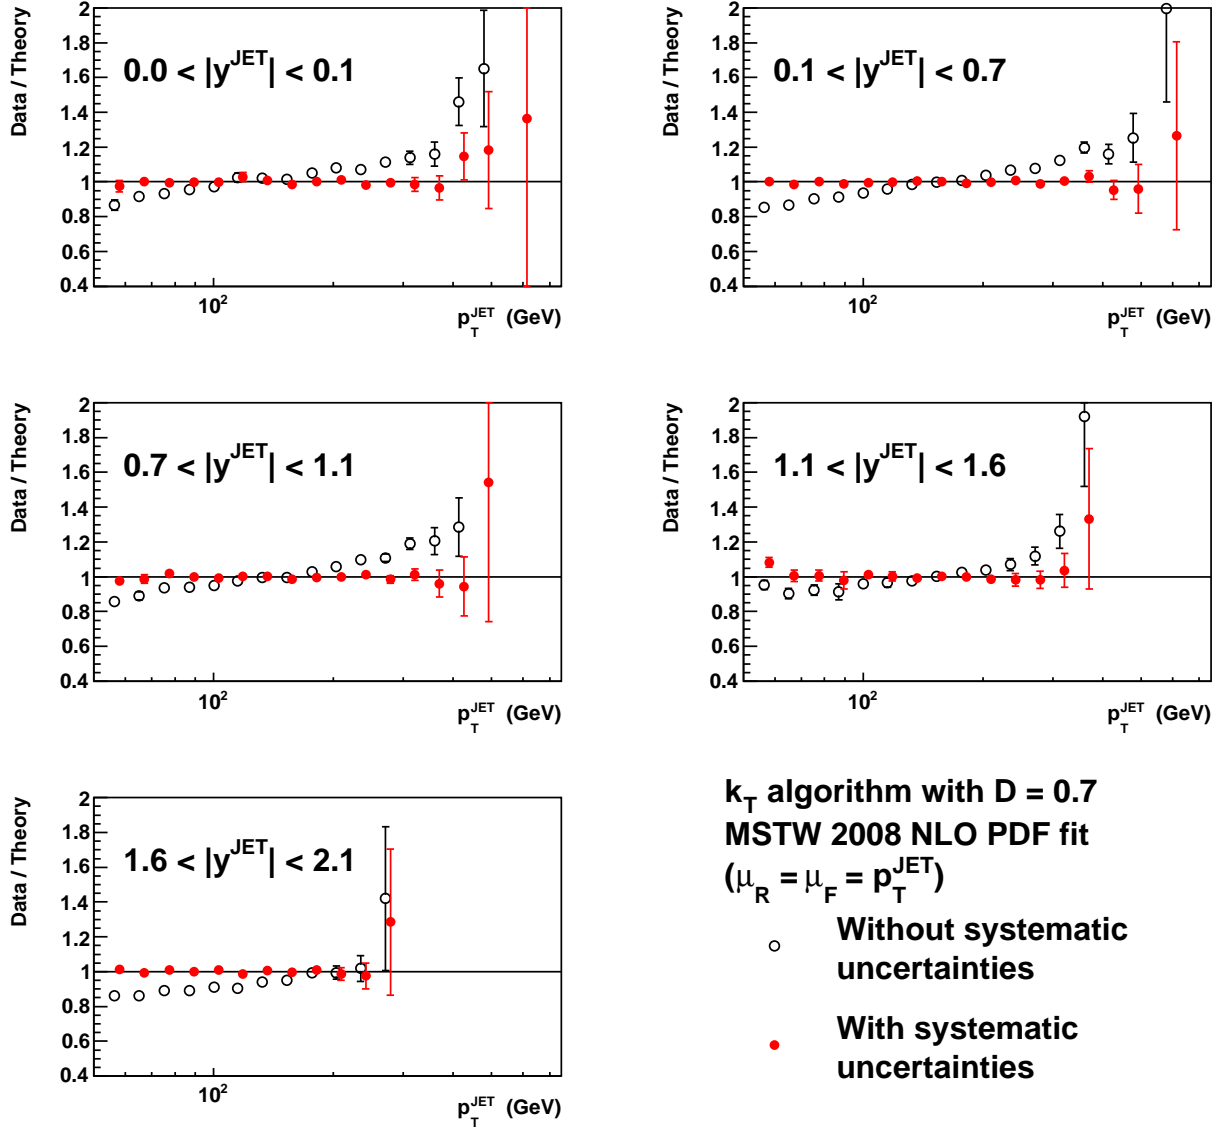


Figure 48: Data/theory ratio for the CDF Run II data obtained using the k_T jet clustering algorithm [54] and the MSTW 2008 NLO fit.

DØ Run II inclusive jet data (cone, $R = 0.7$) **MSTW 2008 NLO PDF fit ($\mu_R = \mu_F = p_T^{\text{JET}}$), $\chi^2 = 114$ for 110 pts.**

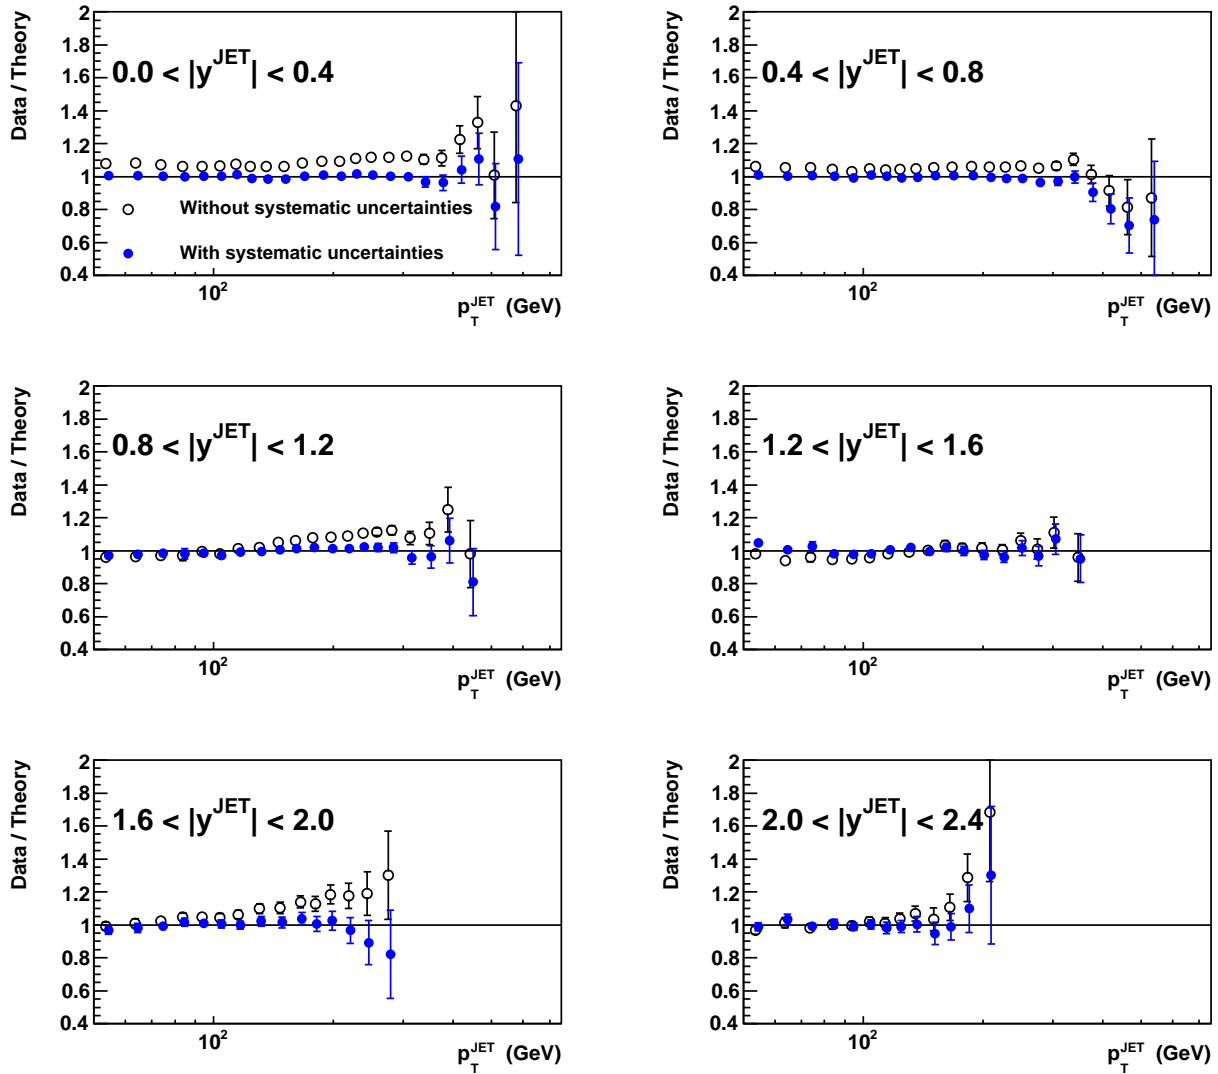


Figure 49: Data/theory ratio for the DØ Run II data [56] and the MSTW 2008 NLO fit.

49. Note that the data points are allowed to shift by the correlated systematic errors according to the χ^2 definition (38). Recall from Table 3 that the CDF Run II inclusive jet data are forced to share the same normalisation as the CDF Run II Z rapidity distribution [52] due to the common luminosity uncertainty. In practice, the shared normalisation of 0.96 at NLO (or 0.99 at NNLO) is determined by the much more constraining Z rapidity distribution.

In Fig. 50 we compare the predictions using the MSTW 2008 NLO PDFs to the CDF Run II data obtained using the Midpoint cone algorithm [55], which were not included in the fit.²⁸ The agreement is generally good, but with the second and third point for $0.7 < |y^{\text{JET}}| < 1.1$, and the first point for $1.1 < |y^{\text{JET}}| < 1.6$, each contributing about 11–13 to the total $\chi^2 = 108$ for 72 points.²⁹ These three points appear simply to be outliers. Replacing the CDF Run II data measured using the k_T jet algorithm by these cone algorithm data and refitting improves the χ^2 by 6 and the χ^2_{global} by 4. The χ^2 of the DØ Run II cone algorithm data is unchanged on refitting. We conclude that there is no significant tension between the three Tevatron Run II data sets on inclusive jet production [54–56] obtained using either the cone [55, 56] or k_T [54] jet clustering algorithms.³⁰

12.3 Impact on the gluon distribution

The impact of the Tevatron Run II jet data on the high- x behaviour of the gluon distribution is interesting. These data prefer a smaller high- x gluon distribution than the Run I data. Indeed, we find that it is no longer necessary to take the input gluon parameterisation (10) in the DIS scheme and then transform to the $\overline{\text{MS}}$ scheme, as was done in the MRST 2004 [18] and MRST 2006 [21] analyses in order to give an enhanced gluon at high x to better fit the Tevatron Run I data. Instead, the input gluon form (10) is now taken directly in the $\overline{\text{MS}}$ scheme.

In Fig. 51 we show the high- x gluon distribution from the MSTW 2008 NLO analysis compared to that obtained from a fit with the Tevatron Run II jet data replaced by Tevatron Run I jet data, and from a fit without any Tevatron jet data. The fit to Tevatron Run I data uses FASTNLO [75] and includes hadronisation corrections that were previously neglected in all global fits, but can reach the level of a $\mathcal{O}(10\%)$ correction in the lower E_T bins. We manage to fit the Tevatron Run I jet data well using the input gluon parameterisation (10) in the $\overline{\text{MS}}$ scheme, without needing to perform the DIS \rightarrow $\overline{\text{MS}}$ transformation used in the MRST 2004 [18] and MRST 2006 [21] analyses. The 68% C.L. uncertainty bands from the fits to Run I and Run II data clearly disagree; the 90% C.L. uncertainty bands also disagree for $x \gtrsim 0.4$ (not shown here).

²⁸We include the complete 24 (excluding luminosity) sources of correlated systematic uncertainty [55]. Note that the data tables originally published in Ref. [55] have received minor corrections, and the corrected data are shown in Fig. 50.

²⁹We note that the rapidity region $0.7 < |y^{\text{JET}}| < 1.1$ corresponds to a crack between the central and plug calorimeters [55].

³⁰If, instead of tying the normalisation to the CDF Run II Z rapidity distribution [52], we allow a separate normalisation for the CDF Run II inclusive jet data, the preferred normalisation at NLO for the Midpoint algorithm data is 1.06, with an improvement in χ^2 to 91 for 72 points. By contrast, the k_T algorithm data prefer a normalisation of 0.98 with almost no improvement in χ^2 compared to the default normalisation of 0.96.

CDF Run II inclusive jet data, $\chi^2 = 108$ for 72 pts.

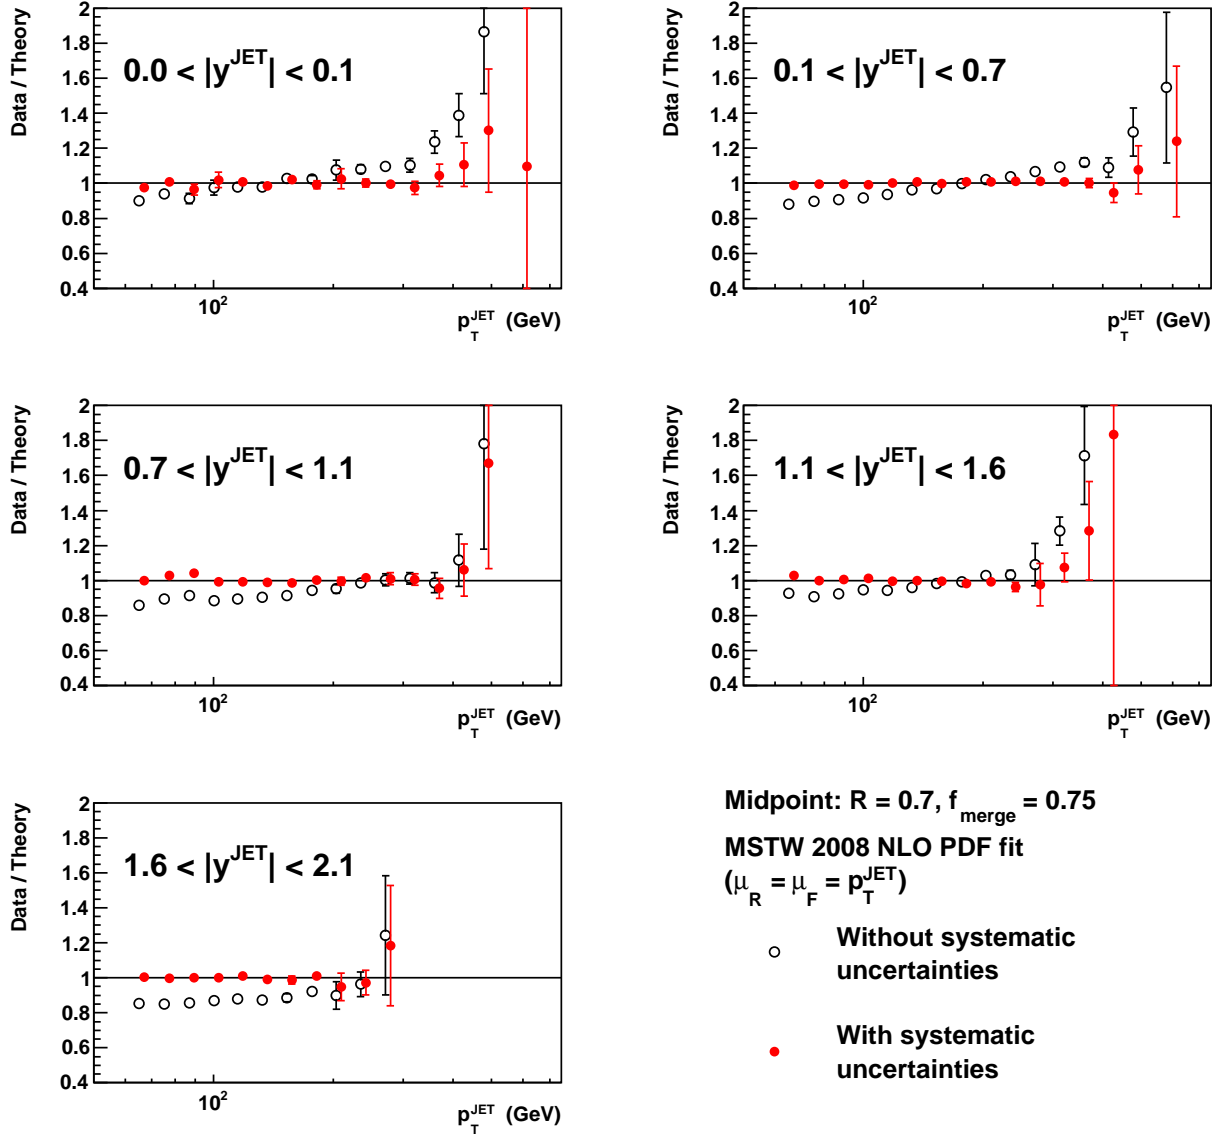



Figure 50: Data/theory ratio for the CDF Run II data obtained using the cone-based Midpoint jet clustering algorithm [55] and the MSTW 2008 NLO fit. These data are not included in the fit.

(a)



(b)




Figure 51: The impact of the Tevatron jet data on the high- x gluon distribution at NLO.

CDFI [179] (33 pts.)	DØI [176] (90 pts.)	CDFII(k_T) [54] (76 pts.)	DØII [56] (110 pts.)	$\Delta\chi^2_{\text{non-jet}}$ (2513 pts.)	$\alpha_S(M_Z^2)$
53	119	64	117	0	0.1197
51	48	132	180	9	0.1214
56	110	56	114	2	0.1202
53	85	68	117	1	0.1204


Table 8: χ^2 values for Tevatron data on inclusive jet production from various global fits. The large bold values indicate the data sets explicitly included in each fit. The “MSTW 2008” fit includes only the Run II data. We also indicate the increase in the χ^2 for all data sets other than Tevatron inclusive jets when these data are added to the fit. Finally, the value of $\alpha_S(M_Z^2)$ is given for each fit. All fits are carried out at NLO with a scale choice of $\mu_R = \mu_F = p_T$.

We also compare in Fig. 51 to the high- x gluon distribution of the ZEUS 2005 Jets analysis [140], which is even smaller than our fit without the Tevatron jet data. This is likely to at least partially be due to the omission of fixed-target data in the ZEUS 2005 Jets analysis. In a global analysis, where $\alpha_S(M_Z^2)$ takes a value of 0.117–0.120, the fixed-target data constrain the high- x gluon distribution to be moderately large, though in general the constraint is correlated with the value of $\alpha_S(M_Z^2)$. The ZEUS 2005 Jets analysis also has a more restrictive input gluon parameterisation than our distribution. Note that although the ZEUS fit [140] included additional data on dijets in photoproduction, which are not included in our fit, these data were found only to reduce the uncertainties without changing the central values, so do not provide an explanation for the discrepancy.

In Table 8 we show the χ^2 values for the Tevatron inclusive jet production data from various global fits. The four lines in the table are fits (i) without Tevatron jet data, (ii) with only Run I data, (iii) with only Run II data, and (iv) with both Run I and Run II data included. The CDF Run I data [179] exist only for a central rapidity bin, and hence are less sensitive to the gluon distribution than the other data sets in Table 8. The χ^2 per data point when fitting only Run I data is significantly larger than 1 for CDF, suggesting unusually large fluctuations, and significantly smaller than 1 for DØ, suggesting an overestimation of errors. The χ^2 per data point when fitting only Run II data is much closer to 1 for both CDF and DØ, suggesting that the Run II data are more reliable than the Run I data. The tension between the Run I and Run II data is evident from Table 8, i.e. the fit to only Run I data gives a poor description of Run II data, and vice versa. The global χ^2 for the fit without Tevatron jet data is 2371 for 2513 points. When including the Run I data it is 9 worse for the same data points, but when including the Run II data it is only 2 worse. The α_S values are also slightly different: the fit without Tevatron jets has $\alpha_S(M_Z^2) = 0.1197$ compared to the fit with Run I data which has $\alpha_S(M_Z^2) = 0.1214$ and the standard fit with Run II data which has $\alpha_S(M_Z^2) = 0.1202$. The Run II jet data are therefore slightly more consistent with the other data sets included in the global fit than the Run I jet data, and lead to a significantly lower value of α_S , which influences the gluon distribution.

In Fig. 52(a) we show that the inclusion of Tevatron jet data does reduce the fractional uncertainty on the high- x gluon, albeit not dramatically. This rather disappointing result is

(a)



(b)




Figure 52: (a) The impact of the Tevatron jet data on the fractional uncertainty of the high- x gluon distribution at NLO. (b) The impact on the high- x gluon distribution of replacing CDF Run II jet data obtained with the k_T jet algorithm [54] by data obtained with the Midpoint cone algorithm [55]. Also shown is the effect of taking an alternative scale choice.

due to the large number of sizeable correlated systematic uncertainties on the jet data, which are free to adjust to best fit the data according to the χ^2 definition (38). There is therefore a considerable trade-off between these systematic shifts and the input gluon parameters.

The high- x gluon distribution from the fit using CDF Run II data obtained with the Midpoint cone algorithm [55] instead of the k_T algorithm [54] lies slightly below the central value of the MSTW 2008 NLO fit, but is well within the uncertainty band, lying slightly above the gluon from the fit without Tevatron jet data; see Fig. 52(b). (The value of $\alpha_S(M_Z^2) = 0.1196$.) This fact confirms the statement made in Section 12.2 that the two CDF Run II data sets are consistent. In Fig. 52(b) we also show the effect of taking a scale choice of $\mu_R = \mu_F = p_T/2$ rather than our default choice of $\mu_R = \mu_F = p_T$ for Tevatron jets. The lower scale choice results in a smaller high- x gluon distribution, which is within the 90% C.L. uncertainty band (but outside the 68% C.L. uncertainty band, not shown here). The value of $\alpha_S(M_Z^2) = 0.1192$ is also smaller with the lower scale choice. The quality of the description to either of the CDF Run II data sets is similar with either scale choice, but the χ^2 of the DØ Run II data is ≈ 10 worse with the lower scale choice. Therefore, overall, the Tevatron Run II jet data prefer our default scale choice of $\mu_R = \mu_F = p_T$. By contrast, when fitting only to Tevatron Run I jet data, the χ^2 of the DØ Run I data is ≈ 10 *better* with $\mu_R = \mu_F = p_T/2$, i.e. the default scale choice used in previous MRST and CTEQ analyses. Finally, in Fig. 52(b) we also show the high- x gluon distribution obtained from a simultaneous fit to Run I and Run II jet data, i.e. the last row in Table 8. Unsurprisingly, it lies in-between the gluon distributions obtained from the two separate fits to either Run I or Run II jet data alone. Note, however, that all gluon distributions shown in Fig. 52(b) are obtained from fits with slightly different values of α_S , and hence are not really directly comparable.


A comparison with earlier gluon distributions is shown in Fig. 53. The plots show that the inclusion of the Run I Tevatron jet data in the MRST 2004 [18] and CTEQ6.6 [87] analyses led to a significantly harder gluon than the present analysis (MSTW 2008) which is based on the Run II Tevatron jet data. (For the MRST 2004 set, where eigenvector PDF sets are not available, we assume that the fractional PDF uncertainties are the same as for the MRST 2001 set.) In fact, with the Run II data, we are almost back to the “soft” MRST 2001 gluon, which did not describe the Tevatron Run I jet data nearly as well as the MRST 2004 analysis.

As explained in Section 5.1, we include the Tevatron Run II jet data in the NNLO analysis on the basis that the NNLO theoretical correction is very likely to be small and a smooth function of p_T , and to be less than the size of the correlated systematic uncertainties on the data. As can be implied from Figs. 51 and 52, the inclusion of these data does not influence the central value of the high- x gluon distribution at all significantly, but does help constrain the uncertainty.

12.4 Jet production at HERA

In Figs. 54 and 55 we show the comparison of the NLO fit to ZEUS data [57, 58] on inclusive jet production, and in Fig. 56 we show the comparison of the NLO fit to H1 data [59]. Jet data have been included previously by the ZEUS Collaboration in their own PDF analysis [140]. H1 have similarly included inclusive jet data in their determination of diffractive PDFs [182]. However, we find that the HERA jet data have little influence on either the central values or the PDF

(a)



(b)




Figure 53: Comparison of the current high- x gluon distribution at NLO with that from the previous MRST [18] and CTEQ [87] analyses.

ZEUS 96-97 inclusive jet data

MSTW 2008 NLO PDF fit, $\chi^2 = 30$ for 30 pts.

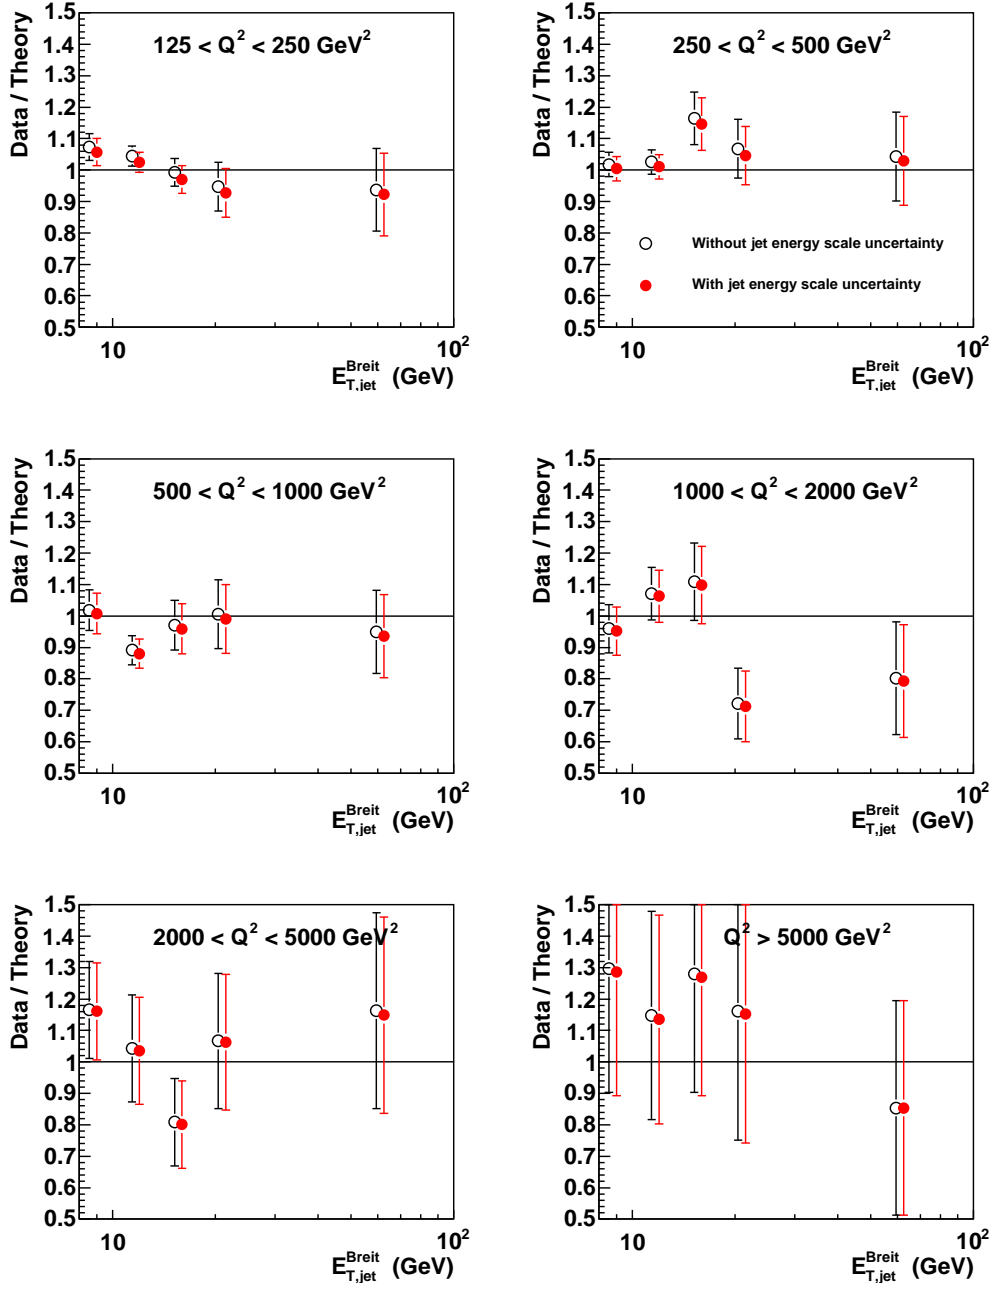


Figure 54: Data/theory ratio for the ZEUS 96–97 inclusive jet data [57] and the MSTW 2008 NLO fit.

ZEUS 98-00 inclusive jet data

MSTW 2008 NLO PDF fit, $\chi^2 = 17$ for 30 pts.

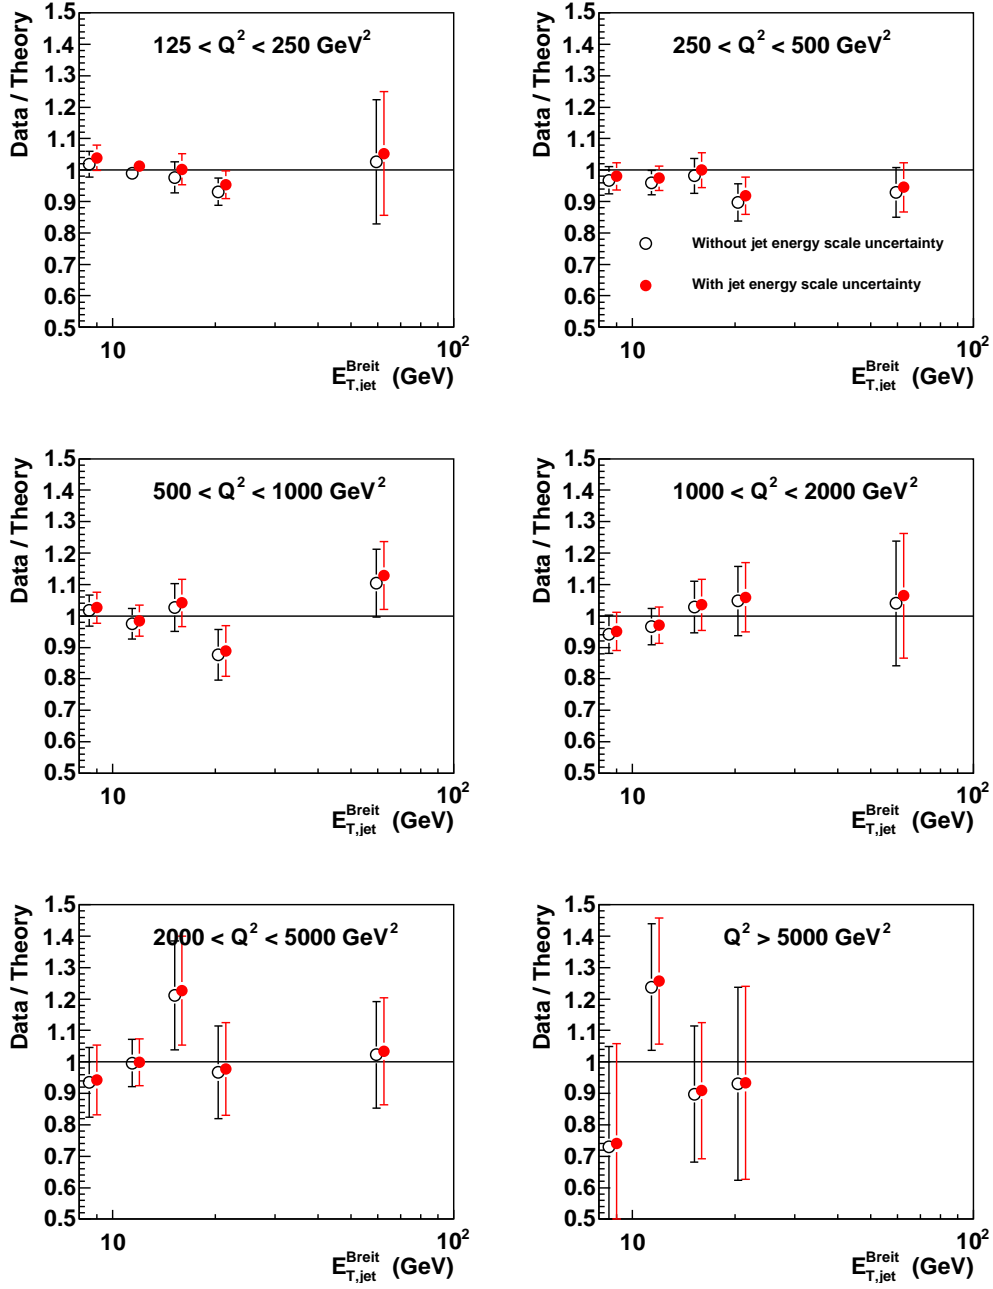


Figure 55: Data/theory ratio for the ZEUS 98-00 inclusive jet data [58] and the MSTW 2008 NLO fit.

H1 99-00 inclusive jet data

MSTW 2008 NLO PDF fit, $\chi^2 = 19$ for 24 pts.

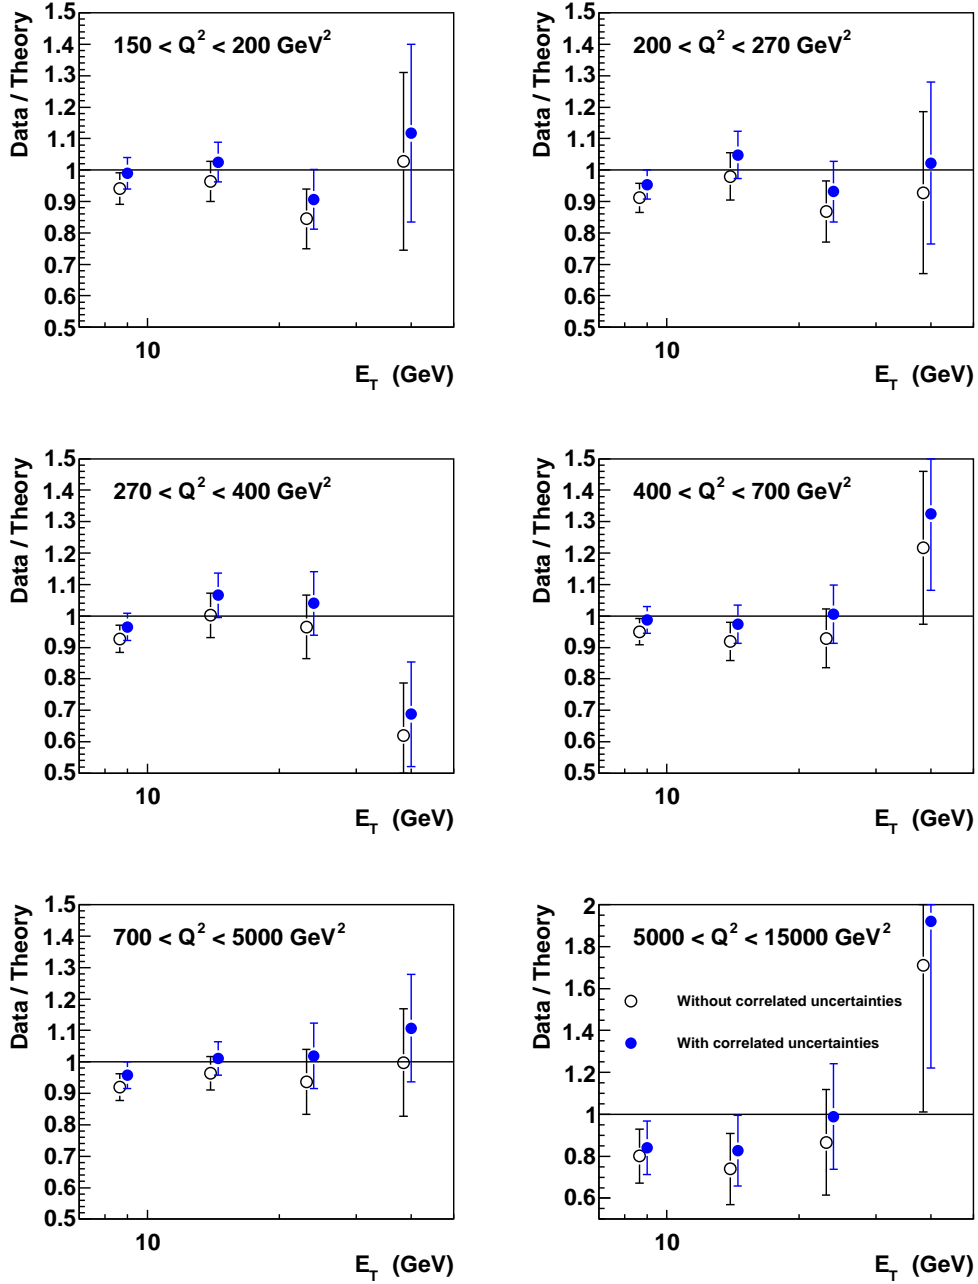


Figure 56: Data/theory ratio for the H1 99–00 inclusive jet data [59] and the MSTW 2008 NLO fit.

uncertainties: the constraint they provide with the current accuracy on the data being mimicked by other data sets in a fully global fit. We do not include the dijet photoproduction data as this involves further systematic uncertainties due to the choice of photon PDFs.

We omit the HERA jet data from the NNLO fit since it is quite probable that the unknown NNLO corrections are fairly large. However, the description of these data is excellent if we use NLO hard cross sections with NNLO parton distributions, i.e. comparable to the fit we obtain at NLO. If the NNLO correction is equivalent to a fairly smooth (and not too large) K -factor, the good fit could likely be maintained by the interplay of data relative to theory using the correlated systematic uncertainties.

13 Low- x gluon distribution and F_L

In this section we outline our results for the small- x gluon distribution and the implications of a measurement of $F_L(x, Q^2)$. This is, in principle, a very important quantity to compare to predictions since it gives an independent test of the gluon distribution at low x to accompany the indirect determination from $\partial F_2(x, Q^2)/\partial \ln Q^2$, which beyond LO results in a perhaps surprisingly small, or even negative, gluon distribution. However, it is also a direct test of the success of extensions of our theory beyond fixed-order calculations. Until recently we have been limited to consistency checks on the relationship between $F_2(x, Q^2)$ and $F_L(x, Q^2)$ at high y , where both contribute to the total cross section $\tilde{\sigma}(x, Q^2) = F_2(x, Q^2) - y^2/(1 + (1 - y)^2)F_L(x, Q^2)$. Extracting F_L data from this requires an extrapolation in y making some theory assumptions. It is more useful to fit directly to $\tilde{\sigma}(x, Q^2)$ where there is a turnover at the highest y . This was discussed in Section 7.1, where we saw that although the full turnover in $\tilde{\sigma}(x, Q^2)$ implied by the data was not seen at NLO, or even in full for the central fit at NNLO, the quality of fit to the data was acceptable given the large uncertainties. Clearly the precision of such studies is limited, and can be affected by systematics, e.g. the photoproduction background uncertainty. However, using the final low-energy run at HERA, a direct measurement of $F_L(x, Q^2)$ at HERA has been possible, and some results have already been presented [183, 184].

The prediction for $F_L(x, Q^2)$ is mainly determined by the form of the gluon distribution extracted from the fit. This is shown in Fig. 57. There is poor stability as one changes the perturbative order of the calculation, particularly at small x and Q^2 . At small x there is a large order-by-order change in the splitting functions, particularly P_{qg} , and this leads to a large variation in the gluon. At NLO a small- x divergence appears for the first time, making the evolution of the quarks much quicker, and hence the NLO gluon is dramatically smaller at small x . Most extractions find a NLO gluon which is small and perhaps valence-like at $Q^2 \lesssim 2 \text{ GeV}^2$ and very small x . However, the gluon parameterisation we have used since Ref. [14] allows the possibility for the gluon distribution to become negative, and for the best fit it chooses to do so for $x \lesssim 10^{-4}$ at $Q^2 = 2 \text{ GeV}^2$. It is frequently claimed that higher twists, in particular gluon recombination effects, will be a cure for this, but simply applying these corrections, on top of NLO leading-twist QCD, results in little change [17, 185]. At NNLO there is an additional $\ln(1/x)/x$ divergence, and at the smallest x the NNLO gluon becomes smaller to compensate, or in practice more negative for very small values of x . However, it is important to note that the gluon beyond leading order does not have a simple *probabilistic* interpretation. This is particularly the case at high orders at small x , where perturbative corrections to cross sections can be large, and even more so in unphysically motivated schemes such as the $\overline{\text{MS}}$ factorisation scheme.

The fact that the gluon is not directly physical is well-illustrated by the NNLO $\mathcal{O}(\alpha_s^3)$ longitudinal coefficient function $C_{L,g}^{(3)}(x)$ [64, 65]. This has a large positive contribution at small x , and this counters the decrease in the small- x gluon. The predictions for $F_L(x, Q^2)$ at LO, NLO and NNLO are shown in Fig. 58. The $F_L(x, Q^2)$ prediction is more stable in going from NLO to NNLO than the gluon distribution at small x . Despite the behaviour of the gluon distribution, the central value of F_L is positive at NNLO for all $x > 10^{-5}$ at $Q^2 = 2 \text{ GeV}^2$, and is greater than at NLO at the smallest x . The uncertainty becomes enormous as x decreases below 10^{-4} ,

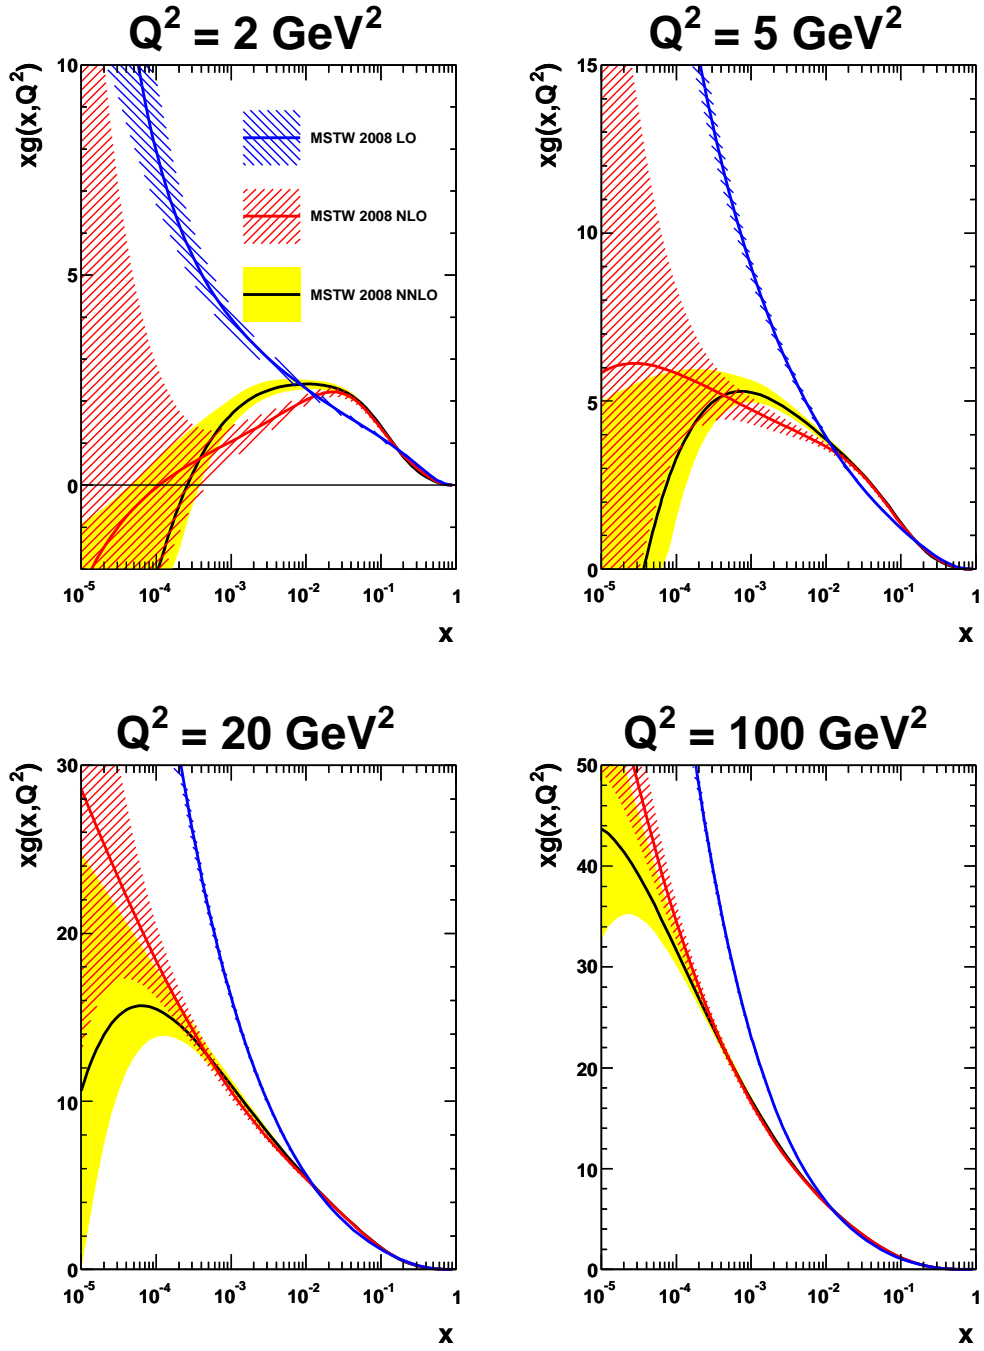


Figure 57: The gluon distribution at LO, NLO and NNLO including the one-sigma PDF uncertainty bands.

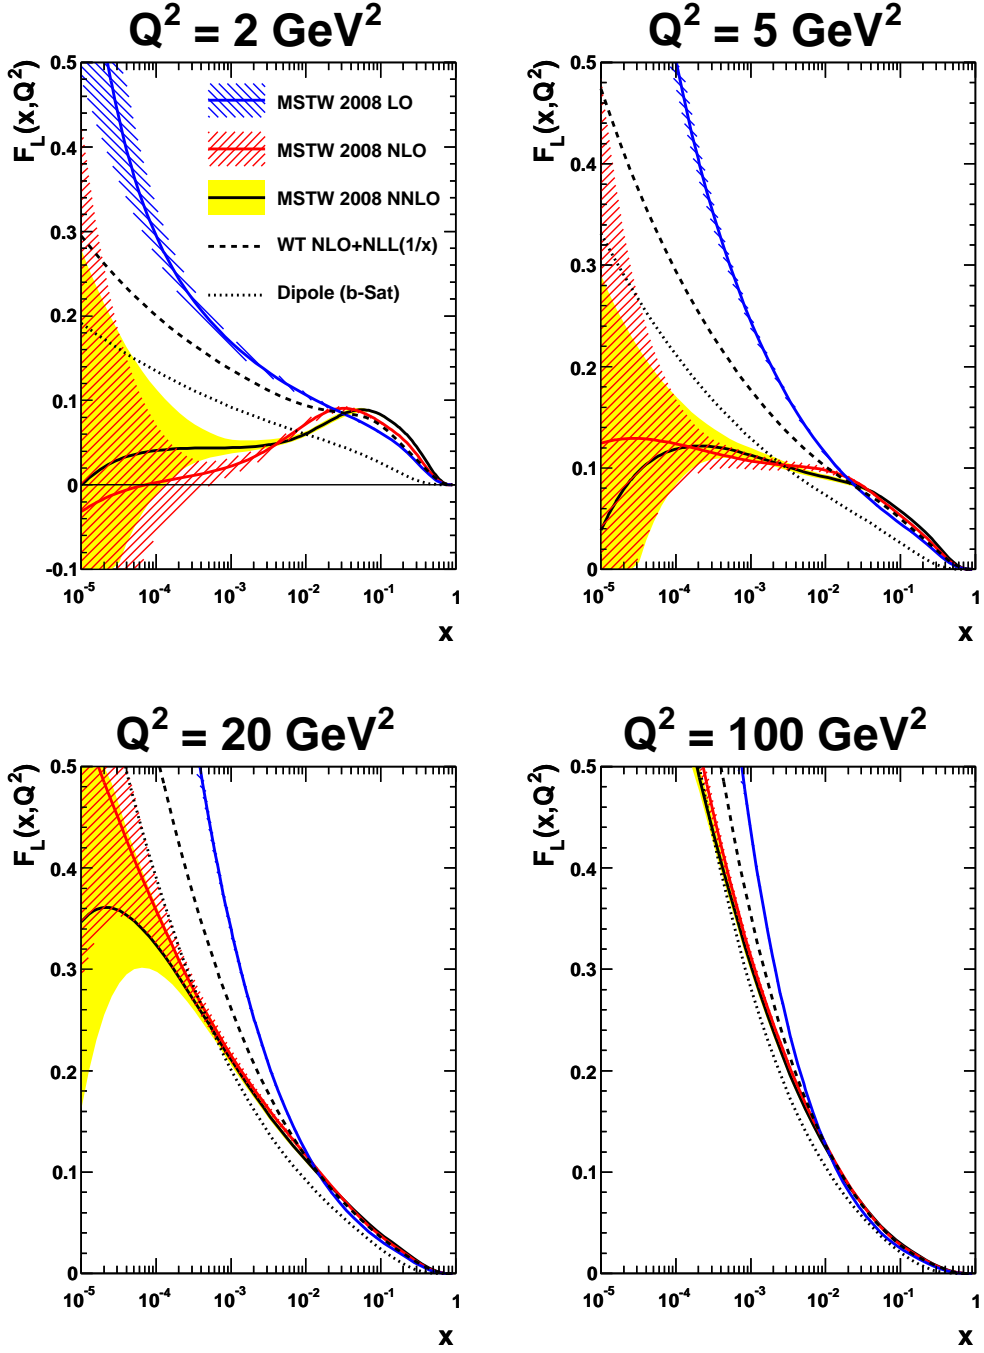


Figure 58: Predictions for the longitudinal proton structure function, $F_L(x, Q^2)$, at LO, NLO and NNLO including the one-sigma PDF uncertainty bands. Also shown are resummed NLL BFKL predictions [144] and dipole model predictions [186].

but at the lowest Q^2 the NLO and NNLO predictions are discrepant in some regions. The LO prediction is far larger than either, reflecting the huge correction in the small- x gluon going to NLO. In order to more clearly separate the relative effect of different gluon distributions, in Fig. 59 we show the predictions for F_L with LO, NLO and NNLO coefficient functions but using common NNLO PDFs. There is a positive correction at high x at each order. NLO then results in a negative correction at lower x , whereas NNLO gives an additional negative correction at intermediate x before becoming positive again at sufficiently small x . However, the value of x at which these transitions occur is very sensitive to Q^2 . As the gluon becomes steeper the regions of the gluon distribution probed in the convolution become more local to the value of x , and small- x divergences in the coefficient functions become less important, along with the relative importance of higher orders decreasing as the coupling becomes weaker. Hence, the ultimate positive effect of the $\ln(1/x)/x$ term in the NNLO coefficient function is only seen to increase the NNLO result above LO for $x > 10^{-5}$ in the two lowest Q^2 bins.

As the relative lack of stability in fixed-order predictions implies, there are various potentially large corrections beyond fixed-order perturbation theory. It is possible that there is a large higher-twist contribution from renormalons in the quark sector [187], and the implications are discussed in Refs. [28, 188]. Perhaps more significantly, since the small- x NNLO correction is itself rather large, even higher orders might be important. There are leading $\ln(1/x)$ terms of the form:

$$xP_{gg}(x) \sim \alpha_S^n \ln^{n-1}(1/x), \quad xP_{qg}(x) \sim \alpha_S^n \ln^{n-2}(1/x) \quad \text{and} \quad xC_{L,g}(x) \sim \alpha_S^n \ln^{n-2}(1/x). \quad (97)$$

A fit which performs a double resummation of leading and next-to-leading $\ln(1/x)$ and running coupling contributions leads to a better fit to small- x data than a conventional perturbative fit [144]. The gluon distribution from this resummed fit, defined in a more physical scheme, is larger at small x and Q^2 than NLO or NNLO, and indeed is always positive for $Q^2 \geq 1 \text{ GeV}^2$. This is reflected also in the prediction for $F_L(x, Q^2)$. Similar approaches [189, 190] all lead to rather comparable results for the calculated splitting functions, but only in Ref. [144] have detailed phenomenological studies taken place. A partially overlapping set of additional corrections are considered in the dipole picture. As with small- x resummations, this approach can be cast in the language of $f(x, k^2)$, the unintegrated gluon distribution, which is directly related to the dipole-proton cross section. The structure functions are obtained by convoluting this dipole cross section with the wave functions for the photon to fluctuate into a quark-antiquark pair. Hence, this picture includes some of the resummation effects, but also higher-twist contributions, and is designed to approach $Q^2 = 0$ smoothly. However, it misses quark and higher- x contributions. In this framework, higher-twist corrections are not small in either F_L or F_T separately, but largely cancel in $F_2 = F_L + F_T$ [143, 191]. Overall the $F_L(x, Q^2)$ predicted in the dipole model approach is steeper at small x than fixed-order predictions, and is automatically stable at lowest Q^2 . The general features are rather insensitive to whether saturation effects are included in the dipole cross section. Resummed NLL BFKL predictions [144] and dipole model predictions [186] are additionally shown in Fig. 58 illustrating these features.

The first published direct measurement of $F_L(x, Q^2)$ can be found in Ref. [183] for $12 \leq Q^2 \leq 90 \text{ GeV}^2$, based on data taken in the last few months of HERA running, when the proton beam

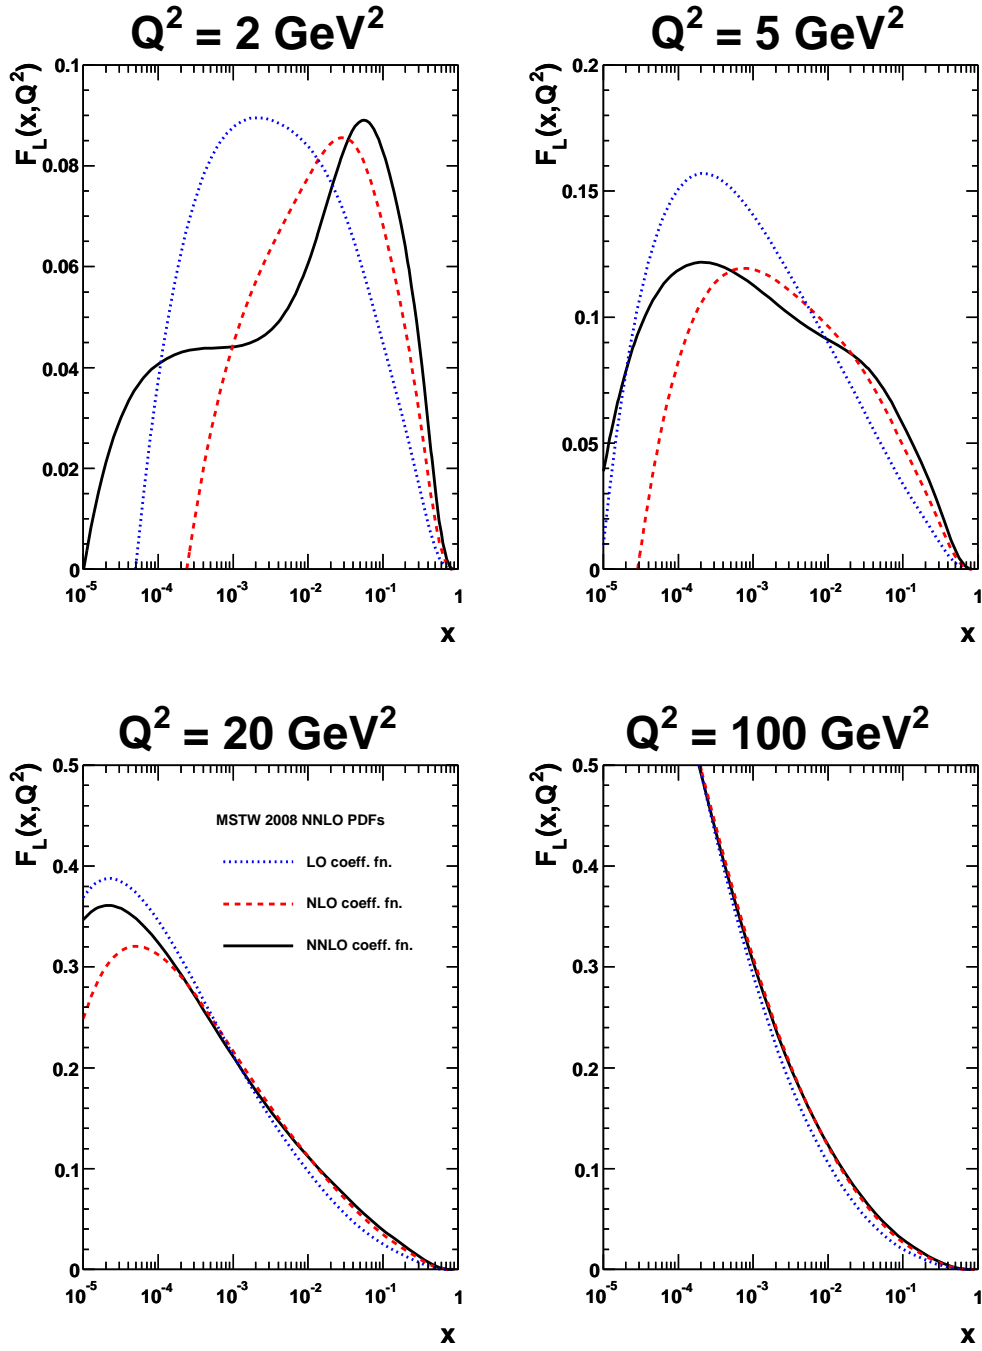


Figure 59: Predictions for the longitudinal proton structure function, $F_L(x, Q^2)$, using NNLO PDFs and LO, NLO and NNLO coefficient functions.




Figure 60: The longitudinal proton structure function, $F_L(x, Q^2)$, measured by H1 [183] compared to predictions at LO, NLO and NNLO including the one-sigma PDF uncertainty bands. Also shown are resummed NLL BFKL predictions [144] and dipole model predictions [186]. (The dependence on the form of the dipole cross section is mild [192].)

energy was lowered from the nominal value of 920 GeV to values of 460 GeV and 575 GeV. In Fig. 60 we show these data compared to our predictions calculated at LO, NLO and NNLO using the appropriate parton distributions at each order. The uncertainty bands obtained from the 40 alternative eigenvector PDF sets are shown at each order. In the region of the existing H1 data the NLO and NNLO predictions agree well with the data, while the LO prediction overshoots the data at the lower Q^2 values. We also show in Fig. 60 the predictions at lower $Q^2 \lesssim 10 \text{ GeV}^2$. We extrapolate the values of x accordingly, by fitting the x values of the H1 data as a power law function of Q^2 , obtaining $x = (1.09 \times 10^{-5})(Q^2/Q_0^2)^{1.28}$ with $Q_0^2 = 1 \text{ GeV}^2$. The NLO and NNLO predictions are very similar along this line, but as seen in Fig. 58 this is partially accidental. Note that the NNLO uncertainty band at low x and Q^2 is smaller than the NLO uncertainty band due to the lesser sensitivity to small- x PDFs in the convolution integrals at NNLO compared to NLO, because the more divergent coefficient function samples the PDFs further from x . Again we also show resummed NLL BFKL predictions [144] and dipole model predictions [186], and both agree well with the published data. For $Q^2 \in [5, 10] \text{ GeV}^2$ the uncertainty in NLO/NNLO predictions for $F_L(x, Q^2)$ due to the gluon uncertainty increases to more than 20%. A good measurement of $F_L(x, Q^2)$ here would automatically improve the gluon determination. Resummations and dipole models suggest a higher low- Q^2 $F_L(x, Q^2)$ by an absolute value of up to 0.15 — well outside the fixed-order uncertainties. So, moreover, a good measurement of $F_L(x, Q^2)$ in this region would start to discriminate between theories.

14 Comparison with other PDF sets

In this section we make a comparison with various other PDF sets obtained from analyses of a variety of data sets by different groups. Some of the most important features pertaining to previous MRST sets have already been mentioned, but here we will present a general overview of the similarities to, and differences from, other sets, and suggest reasons to explain these differences.

We begin with a historical perspective. In Fig. 61 we show the current NNLO PDFs compared to those from MRST 2006 [21]. The comparison would in general terms be rather similar if comparing our current NLO set to the (unpublished) 2006 NLO set, but because the change in $\alpha_S(M_Z^2)$ has been slightly larger at NNLO some differences in PDFs are a little more apparent in the NNLO comparison. Comparing with older NLO sets would mean a change in the definition of the GM-VFNS, and the consequences of this have already been illustrated in Fig. 3. Beginning from the top of Fig. 61, we see that the shape of both valence distributions has changed significantly, though there is still compatibility with MRST 2006 within the 90% C.L. uncertainty bands. This is partially due to the smaller value of $\alpha_S(M_Z^2)$ — the coefficient functions give a positive effect at very high x and negative for $x \sim 0.1$, and this is smaller in 2008 than 2006, so the quarks compensate. However, the details, and the large changes in d_v are due to the new Tevatron Run II W and Z data and the NuTeV and CHORUS neutrino DIS data. For $x < 0.01$ the form of both valence quark distributions is determined mainly by the requirement to satisfy the number sum rules. The gluon distribution has changed significantly: the smallness at high- x being mainly due to the Tevatron Run II jet data, as discussed in detail in Section 12. This is part of the reason for the decrease in $\alpha_S(M_Z^2)$, and the weaker coupling and the momentum sum rule both contribute to the significantly larger gluon distribution at lower x in the 2008 set. This feeds through, via evolution, to the sea quarks being higher at small x , though well within previous uncertainties. For the gluon and light quarks the uncertainty is generally a little larger than for the MRST 2006 set. This is due largely to the inclusion of data normalisation uncertainties into the PDF error analysis, but also due to an improvement in the flexibility of the parameterisation for the gluon and down valence distribution, and for light sea quarks the independent determination of the strange distribution also leads to an increase. The fact that we fit the strange distribution directly results in the very large increase in the uncertainty for $x > 0.001$. But as we assume that the input strange has the same shape as $x \rightarrow 0$ as the light quark distributions, the uncertainty becomes the same as for u and d at very small x and high Q^2 .

In Fig. 62 we show the NLO PDFs at $Q^2 = 10^4 \text{ GeV}^2$ compared to those from CTEQ6.6 [87]. The CTEQ6.6 analysis uses a very similar total set of data, but is not as up-to-date, in particular not including any Tevatron Run II data or the neutrino DIS data sets from NuTeV and CHORUS. Overall, there is good agreement between the two sets, with very few regions where the 90% C.L. limits do not overlap. It is striking that the difference between the MSTW 2008 and CTEQ6.6 valence quarks is rather similar to that between the MSTW 2008 and MRST 2006 distributions. This suggests that the difference is largely due to the new NuTeV, CHORUS and Tevatron Run II data included in our analysis. The strange quark distribution is rather different, though the 90% C.L. uncertainty bands overlap. This is due to the different assumptions made

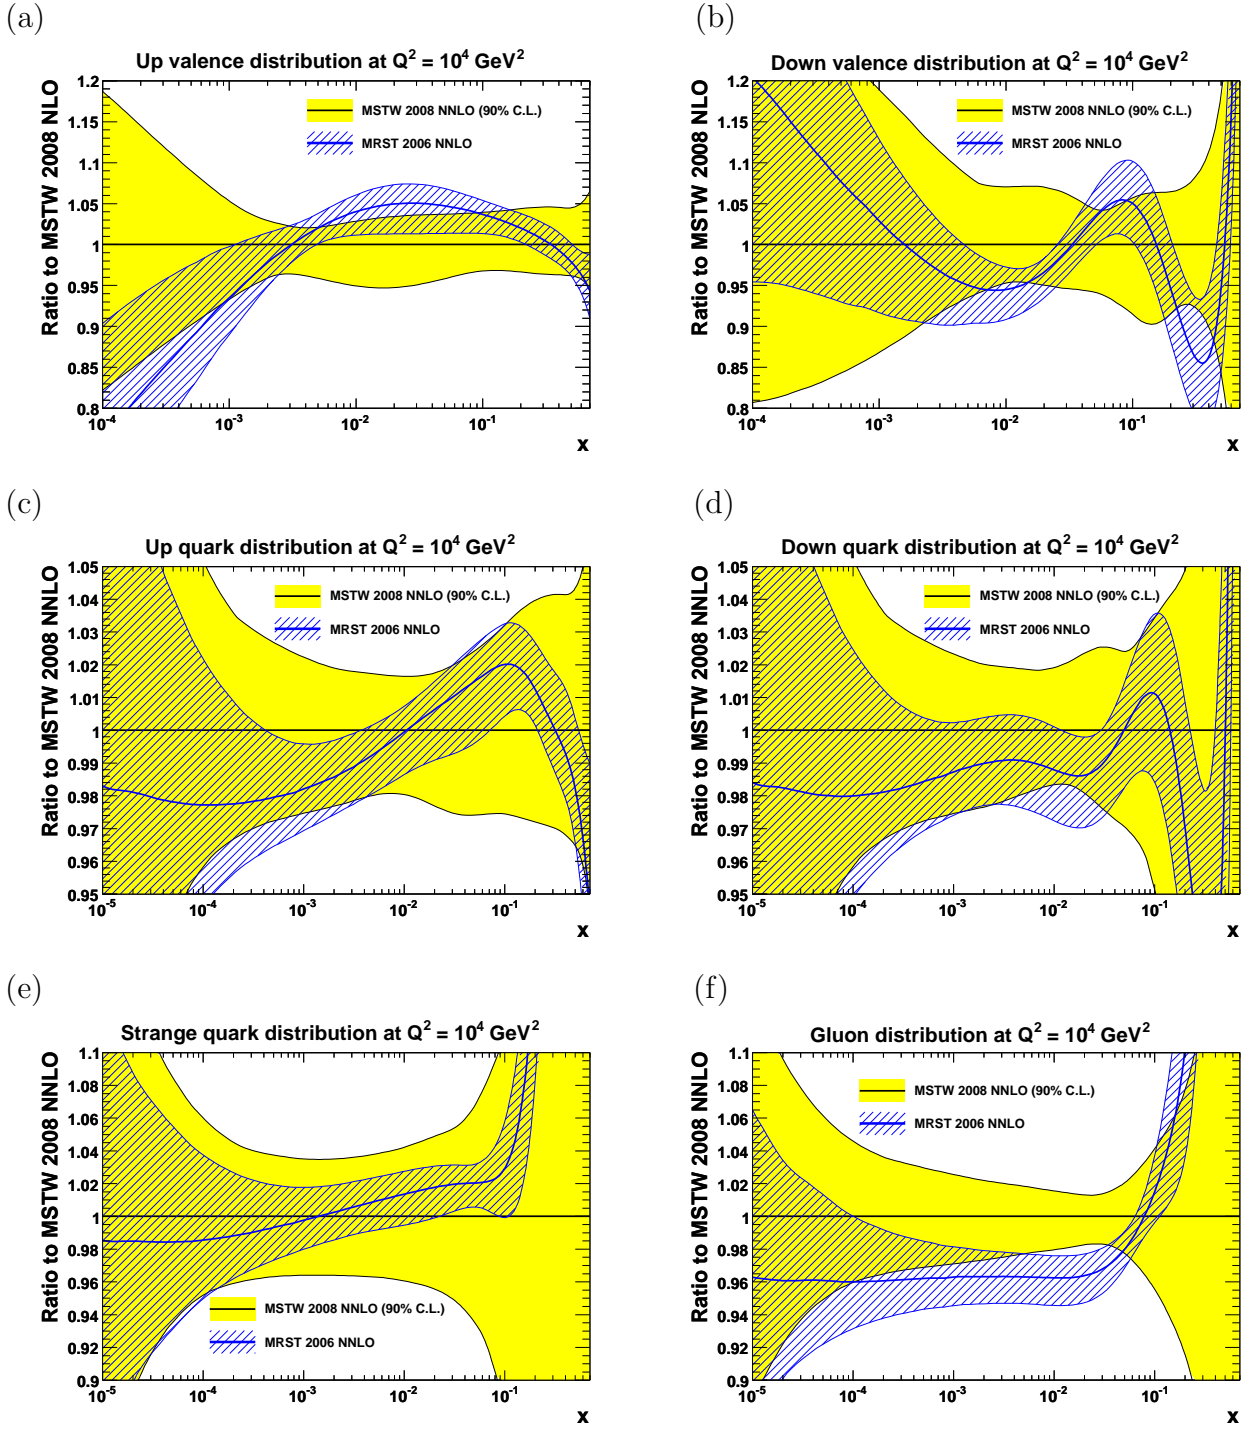


Figure 61: MSTW 2008 NNLO PDFs compared to MRST 2006 NNLO PDFs.

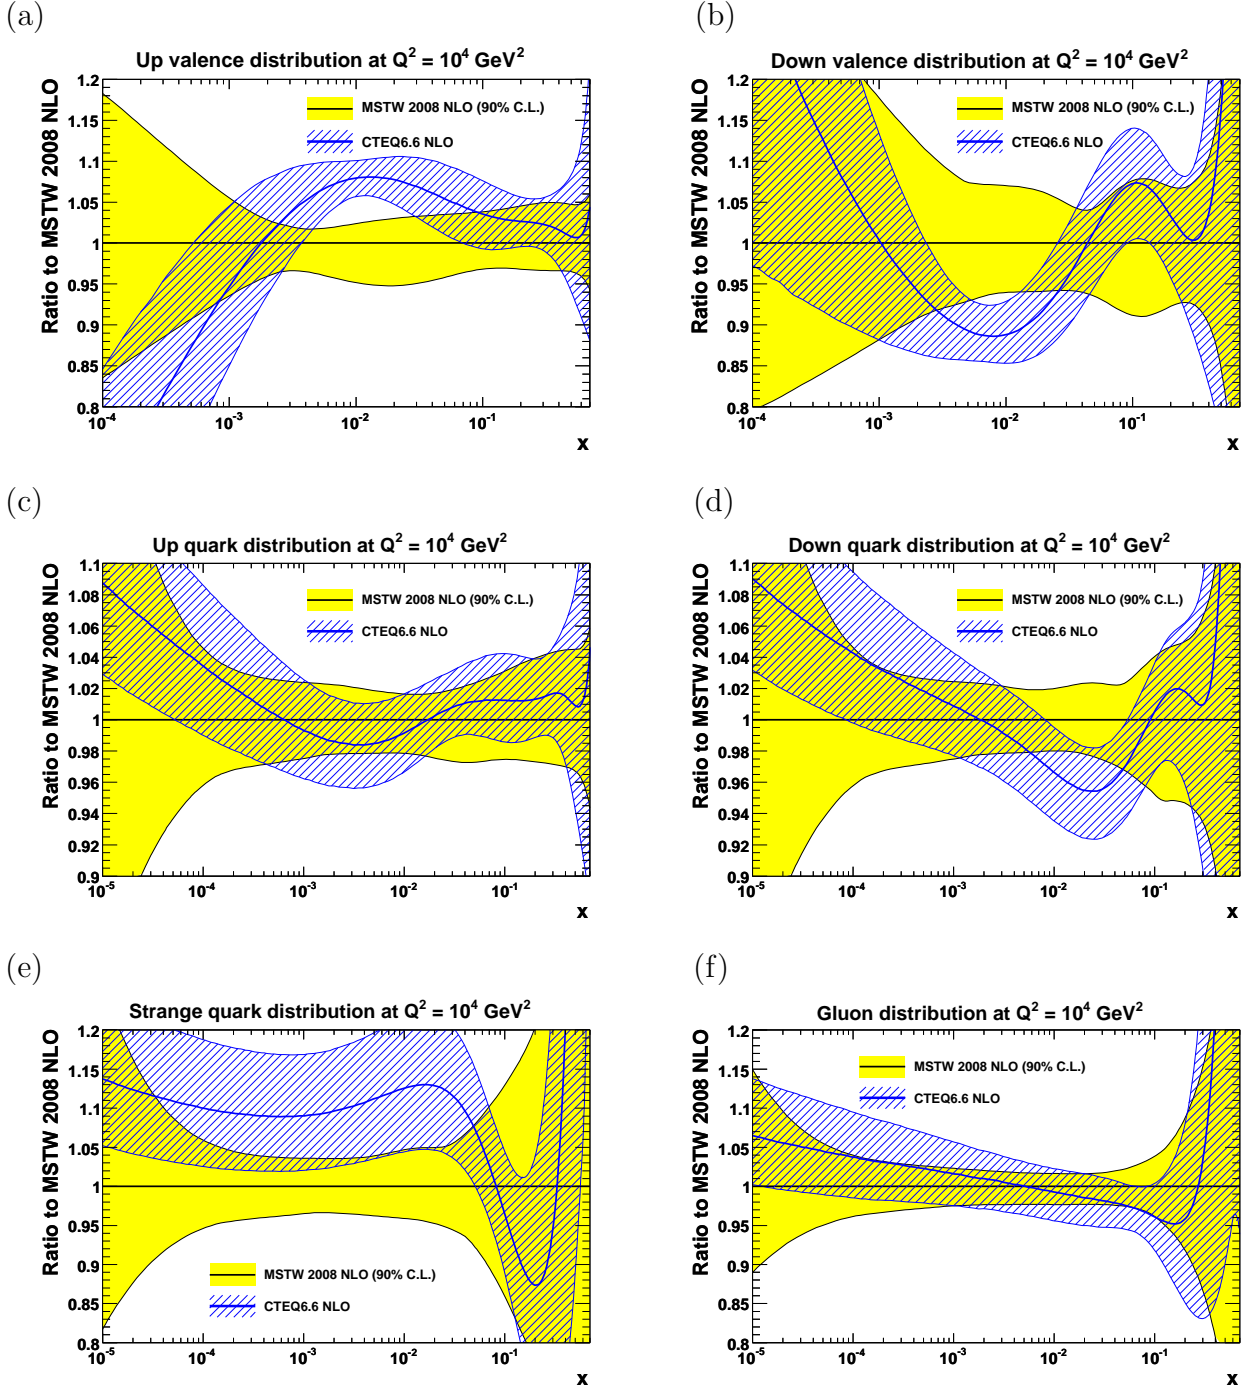


Figure 62: MSTW 2008 NLO PDFs compared to CTEQ6.6 NLO PDFs.

in parameterising the strange distribution, and was addressed in Section 6.5. As in previous comparisons the gluons have a systematic difference. Since MSTW use the flexible parameterisation which allows the distribution to be negative at very small x at input (and at 1.69 GeV^2 , the CTEQ input scale) while the CTEQ input is valencelike, the latter is bigger at the smallest x values. The gluon in the CTEQ6.6 set is determined at high x by Run I Tevatron jet data, so is higher in this region, and from the momentum sum rule must be smaller at intermediate $x \sim 0.1$ to compensate. However, the greater similarity in the heavy flavour treatments now employed perhaps leads to better agreement in the intermediate x range than previously. The comparison of the u and d distributions for $x < 0.1$, i.e. outside the valence region, very much reflects the gluon difference. Finally, we note that the MSTW 2008 uncertainty bands are not as small compared to CTEQ bands as in previous comparisons, despite the addition of more data and on average a smaller tolerance for the former. This is a reflection of the data normalisation uncertainties and the improved flexibility in parameterisations.

In Fig. 63 we show the NLO PDFs compared to those from Alekhin [139]. As one can see there are many significant differences. The Alekhin NLO PDFs are obtained from a fit to structure function data only. An update including some Drell–Yan data, and in particular the E866/NuSea pd/pp data, was made in Ref. [168], but only NNLO sets were made available. Since NLO sets are more widely used, and there are some systematic differences in the approaches at NNLO, we instead compare to the NLO set. There is clearly a significant discrepancy in quark distributions everywhere, despite the fact that largely the same neutral-current DIS data are fit. This highlights the importance of fitting neutrino DIS data and Drell–Yan and vector boson hadroproduction data to obtain the quark flavour decomposition. The absence of Tevatron jet data plays a part in the gluon distribution in the Alekhin set being much smaller at $x > 0.3$. The smaller value of $\alpha_S(M_Z^2) = 0.117$ also contributes to differences, as does the fact that the PDFs correspond to the ZM-VFNS. However, these certainly do not explain the major systematic differences.

Finally we compare to the very recent NNPDF1.0 set [137], which uses proton and deuterium target DIS data, as well as CHORUS data (without nuclear corrections), again in the ZM-VFNS. This approach relaxes the restriction in uncertainty due to fixed parameterisation limitations, though in practice relies on a very large number of parameters. It also proceeds by making a Monte Carlo sample of the distribution of experimental data by generating a large number of replicas of data centred on each data point with full inclusion of the information from errors and their correlations. Each replica is used to generate a PDF set and the mean for a quantity obtained from averaging, and uncertainties from standard deviations. The data are split into training and validation sets for each replica, and the χ^2 for one is monitored while the other minimised, thus avoiding over-complicating the input PDFs by stopping when the fit to the validation sets stop improving. As with the Alekhin set there are some significant systematic differences between these and the MSTW 2008 sets. However, the use of some neutrino DIS data and the fact that a fixed $\alpha_S(M_Z^2) = 0.119$ is used in NNPDF1.0, together with a more flexible input parameterisation compared to Alekhin, means that there are less dramatic variations. Some of the most significant variations are shown in Fig. 64 where we compare our up valence and up antiquark distribution to those from NNPDF1.0 [137]. Some significant differences are also observed in the total up quark distribution and the small- x gluon at high Q^2 . From Fig. 64

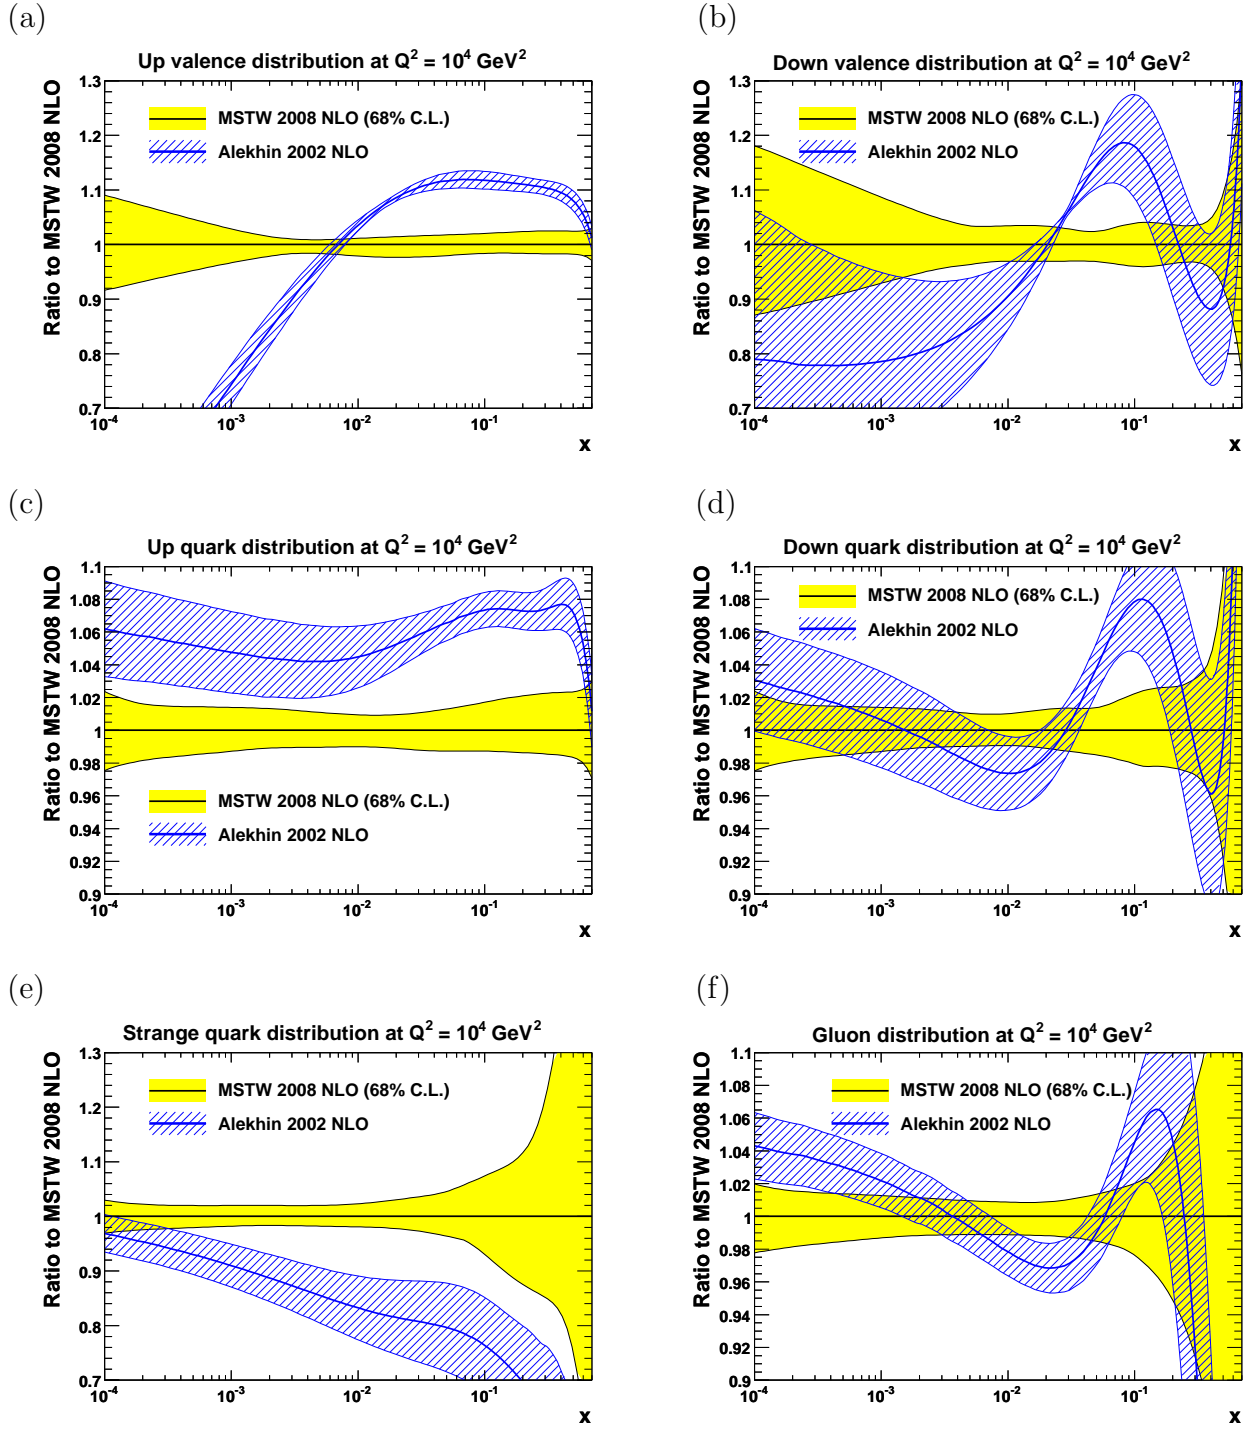


Figure 63: MSTW 2008 NLO PDFs compared to Alekhin NLO PDFs.

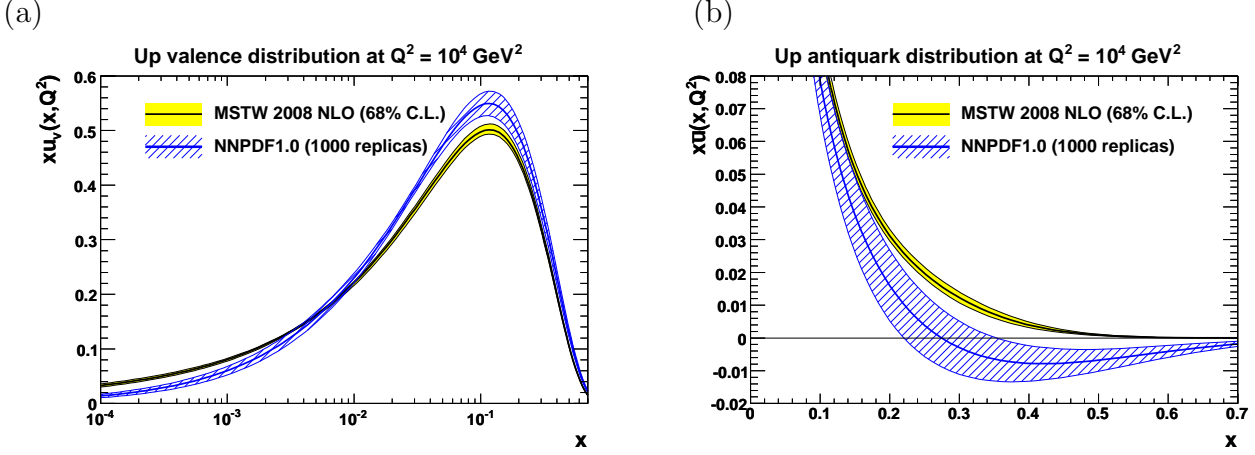


Figure 64: MSTW 2008 NLO (a) up valence quark and (b) high- x \bar{u} distribution compared to those from NNPDF1.0.

we see there are sizeable differences in the central values of the up valence distribution, even where the sea is very small, perhaps due to the fact that relative data set normalisations are fitted in our approach but not in the NNPDF1.0 analysis. The difference in the valence distribution is also undoubtedly due to the lack of data distinguishing between valence and sea quarks. Indeed, the most striking disagreement is in the up antiquark distribution at very large x . In the global fit this would be constrained to be positive by Drell–Yan data. Some differences in central values is not a great surprise but it is more surprising that the error propagation used in the NNPDF1.0 analysis does not completely encompass the MSTW 2008 values. Perhaps this is a reflection of the fact that some data sets used in the fit are incompatible, and discrepancies will sometimes occur, as we found comparing our benchmark fit and standard fit for the high- x gluon in Section 6.6. Since it is the uncertainties that are currently most interesting for the NNPDF1.0 set we concentrate our comparison on this. In Fig. 65 we compare our fractional uncertainties to those from NNPDF1.0 [137]. For total quark and gluon distributions the NNPDF uncertainties are only a little larger, and it must be remembered that we have many more data sets in the fit. For the strange quark the fact that we fit directly, whereas in NNPDF1.0 the strange distribution at input is a fixed fraction of the total light sea, actually leads to our uncertainty being larger for $0.01 \lesssim x \lesssim 0.2$. The increased flexibility in the NNPDF parameterisation does lead to a larger uncertainty at high x and small x for valence distributions, though we can argue that we use more constraining data for $0.01 \lesssim x \lesssim 0.7$, and have less flexibility from the sum rule outside this range. The NNPDF1.0 valence quark uncertainty narrows in the region $x \simeq 0.025$ at $Q^2 = 10^4 \text{ GeV}^2$ (and at $x \simeq 0.06$ at $Q^2 = 2 \text{ GeV}^2$) which we do not understand, as this is just below the lower limit of the BCDMS data, which is one of the best constraints. An update, NNPDF1.1 [159], has recently been reported. The uncertainties are larger in this analysis due both to a randomisation of the preprocessing exponents and due to extra freedom introduced in the strangeness sector. Since there are no dimuon data used to constrain the strange PDF, and very little constraint from any other data set, the uncertainty on $s(x, Q^2) + \bar{s}(x, Q^2)$ in NNPDF1.1

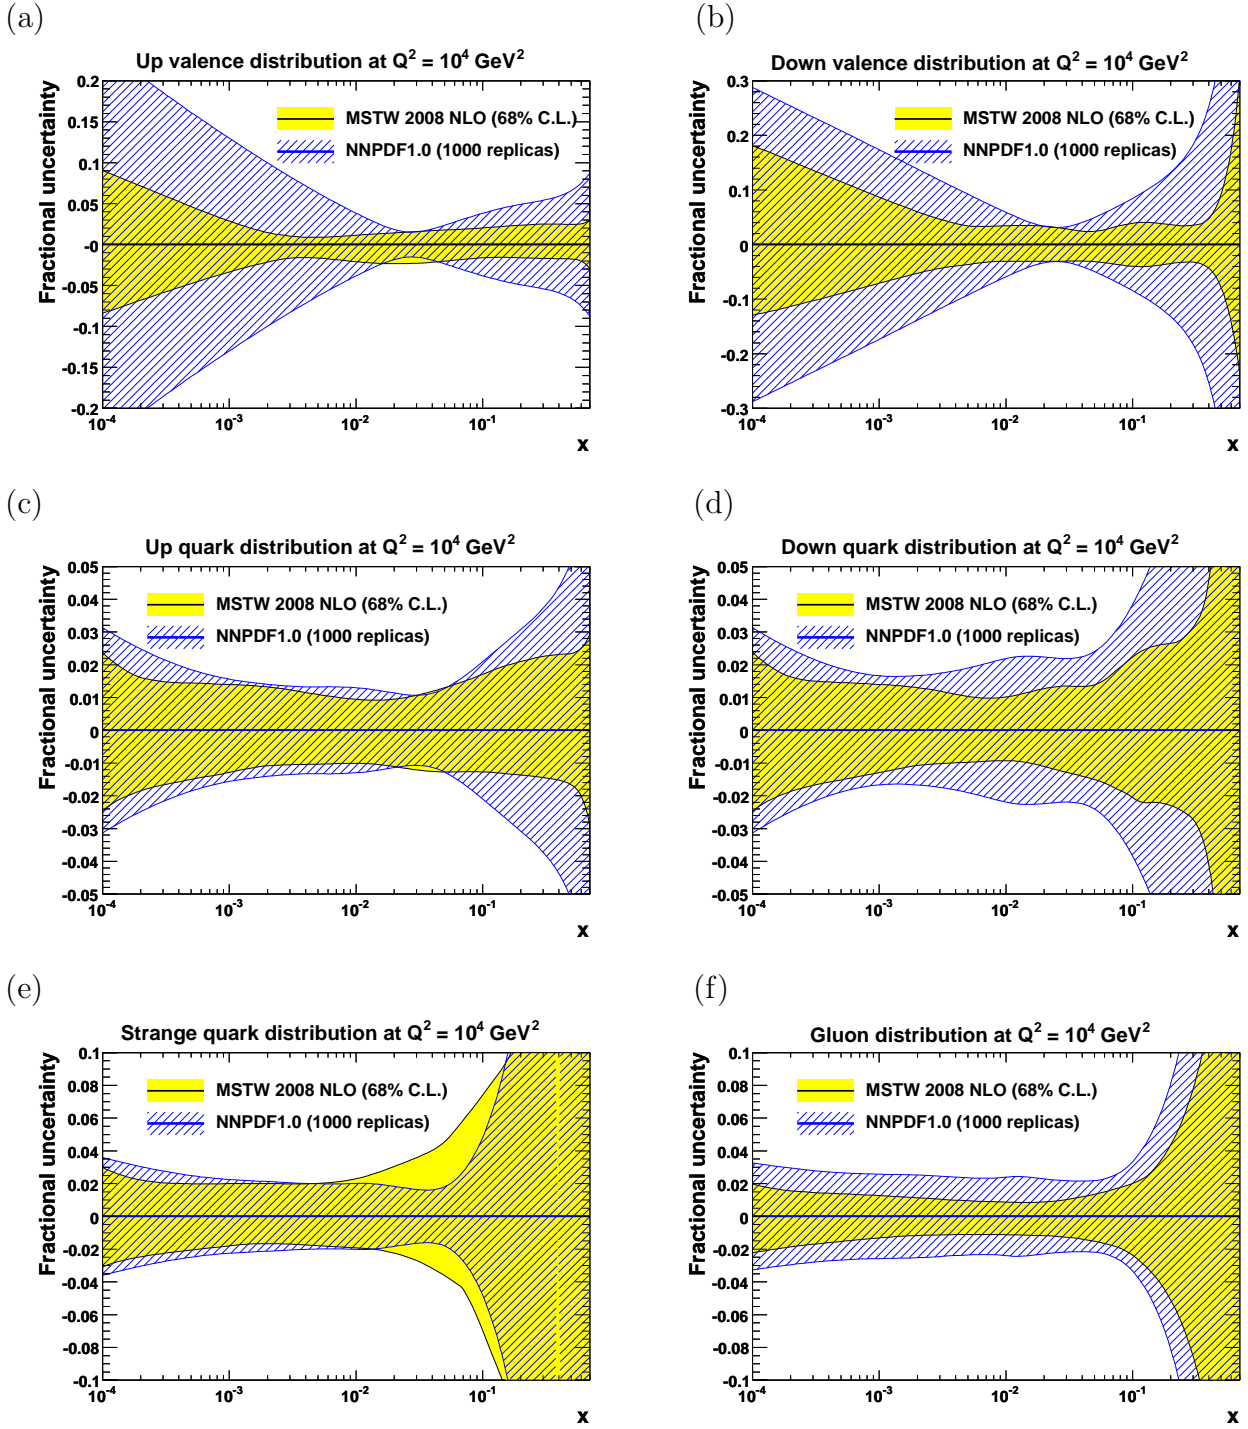


Figure 65: MSTW 2008 NLO PDF uncertainties compared to those from NNPDF1.0.

is huge. This feeds into the other PDFs, with uncertainties increasing by factors of three for many quark combinations, and it is difficult to tell what effect the randomisation alone has. This illustrates how important the direct constraint on the strange PDF, and some theoretical input in the parameterisation, can be when obtaining PDFs with genuinely representative uncertainties.

15 W and Z total cross section predictions

In this section we present and discuss predictions for the total cross sections for W and Z production multiplied by leptonic branching ratios at the Tevatron and LHC. To begin with we consider $W \equiv W^+ + W^-$ production, and subsequently consider the separate W^+ and W^- total cross sections at the LHC. (At the Tevatron, the W^+ and W^- total cross sections are equal.)

To avoid confusion, we first define exactly what we mean by the total W and Z cross sections. We assume a narrow-width approximation, in which the W and Z are treated as on-shell particles, of mass M_W and M_Z respectively, in the calculation of the production cross sections. These cross sections are then multiplied by the appropriate leptonic branching ratios. We assume that the cross sections measured at the leptonic level over a restricted region of phase space can be corrected back to this “theoretical” quantity. We present predictions at LO, NLO and NNLO in perturbative QCD, using a program based on that of van Neerven *et al.* [193] incorporating subsequent corrections from Harlander and Kilgore [194]. We use LO electroweak perturbation theory, with the qqW and qqZ couplings defined by

$$g_W^2 = G_F M_W^2 / \sqrt{2}, \quad g_Z^2 = G_F M_Z^2 \sqrt{2}. \quad (98)$$

The electroweak mixing angle $\sin^2 \theta_W^{\text{eff}}$ determines the relevant strengths of the vector and axial-vector Z couplings to the quarks. Since our cross section predictions depend to some extent on the electroweak parameters, we use the latest (2008) values from the Particle Data Group (PDG) [132]. We sum over u, d, s, c, b quarks in the initial state, using the 2008 PDG (unitarity constrained) values of the CKM weak mixing matrix elements $|V_{qq'}|$ in the case of W^\pm production.

The MSTW 2008 predictions for the total W and Z cross sections, including the one-sigma PDF uncertainties, are given for the Tevatron and LHC in Table 9. It is seen that the NNLO PDF uncertainties are around 2% at both the Tevatron and LHC. The uncertainty coming from neglected higher-order QCD corrections beyond NNLO, estimated by allowing the factorisation and renormalisation scales to vary so that the ratios $\mu_F/M_{W,Z}$, $\mu_R/M_{W,Z}$ and μ_F/μ_R all lie in the range between 0.5 and 2 (as in, for example, Ref. [195]), is around $\pm 0.5\%$ at the Tevatron and $\pm 1\%$ at the LHC. Uncertainties not well accounted for by scale variations, in particular the effect of small- x resummations, variations in GM-VFNS and changes in $\alpha_S(M_Z^2)$ may all lead to further uncertainties. These are difficult to quantify at NNLO, but at the LHC may be 1–2% in each case. For the Tevatron only the last, i.e. $\alpha_S(M_Z^2)$, is likely to be significant. In Fig. 66 we compare the MSTW 2008 NLO and NNLO predictions for the W and Z total cross sections to Tevatron data from CDF [196] and DØ [197]. The agreement is good. The experimental error, which is dominated by the $\sim \pm 6\%$ uncertainty on the machine luminosity measurement, is significantly larger than the uncertainty on the theoretical predictions.

For both theory and experiment, the uncertainty in the W and Z cross section *ratio* is much smaller than on the individual cross sections. This is illustrated in Figs. 67 and 68, which show two-dimensional plots of the MSTW 2008 NLO and NNLO W cross sections against the corresponding Z cross sections, including the error ellipses calculated using (54) and (55) with the one-sigma error sets.³¹ We compare to the central values of cross sections calculated with

³¹As noted in Ref. [87], the error ellipse calculated with e.g. one-sigma PDFs corresponds to a confidence level smaller than one-sigma in two dimensions [131, 132].

Tevatron, $\sqrt{s} = 1.96$ TeV	$B_{\ell\nu} \cdot \sigma_W$ (nb)	$B_{\ell+\ell-} \cdot \sigma_Z$ (nb)	R_{WZ}
MSTW 2008 LO	$1.963^{+0.025}_{-0.028}$ (+1.2%) (-1.4%)	$0.1788^{+0.0023}_{-0.0025}$ (+1.3%) (-1.4%)	$10.98^{+0.02}_{-0.03}$ (+0.2%) (-0.3%)
MSTW 2008 NLO	$2.659^{+0.057}_{-0.045}$ (+2.1%) (-1.7%)	$0.2426^{+0.0054}_{-0.0043}$ (+2.2%) (-1.8%)	$10.96^{+0.03}_{-0.02}$ (+0.3%) (-0.2%)
MSTW 2008 NNLO	$2.747^{+0.049}_{-0.042}$ (+1.8%) (-1.5%)	$0.2507^{+0.0048}_{-0.0041}$ (+1.9%) (-1.6%)	$10.96^{+0.03}_{-0.03}$ (+0.2%) (-0.2%)

LHC, $\sqrt{s} = 10$ TeV	$B_{\ell\nu} \cdot \sigma_W$ (nb)	$B_{\ell+\ell-} \cdot \sigma_Z$ (nb)	R_{WZ}
MSTW 2008 LO	$12.57^{+0.13}_{-0.19}$ (+1.1%) (-1.5%)	$1.163^{+0.011}_{-0.017}$ (+1.0%) (-1.5%)	$10.81^{+0.02}_{-0.02}$ (+0.2%) (-0.2%)
MSTW 2008 NLO	$14.92^{+0.31}_{-0.24}$ (+2.1%) (-1.6%)	$1.390^{+0.029}_{-0.022}$ (+2.1%) (-1.5%)	$10.73^{+0.02}_{-0.02}$ (+0.2%) (-0.2%)
MSTW 2008 NNLO	$15.35^{+0.26}_{-0.25}$ (+1.7%) (-1.6%)	$1.429^{+0.024}_{-0.022}$ (+1.7%) (-1.6%)	$10.74^{+0.02}_{-0.02}$ (+0.2%) (-0.2%)

LHC, $\sqrt{s} = 14$ TeV	$B_{\ell\nu} \cdot \sigma_W$ (nb)	$B_{\ell+\ell-} \cdot \sigma_Z$ (nb)	R_{WZ}
MSTW 2008 LO	$18.51^{+0.22}_{-0.32}$ (+1.2%) (-1.7%)	$1.736^{+0.019}_{-0.028}$ (+1.1%) (-1.6%)	$10.66^{+0.02}_{-0.02}$ (+0.2%) (-0.2%)
MSTW 2008 NLO	$21.17^{+0.42}_{-0.36}$ (+2.0%) (-1.7%)	$2.001^{+0.040}_{-0.032}$ (+2.0%) (-1.6%)	$10.58^{+0.02}_{-0.02}$ (+0.2%) (-0.2%)
MSTW 2008 NNLO	$21.72^{+0.36}_{-0.36}$ (+1.7%) (-1.7%)	$2.051^{+0.035}_{-0.033}$ (+1.7%) (-1.6%)	$10.59^{+0.02}_{-0.03}$ (+0.2%) (-0.3%)

Table 9: Predictions for $W \equiv W^+ + W^-$ and Z total cross sections at the Tevatron and LHC, including the one-sigma PDF uncertainties, and their ratio R_{WZ} . We take $\mu_R = \mu_F = M_{W,Z}$.

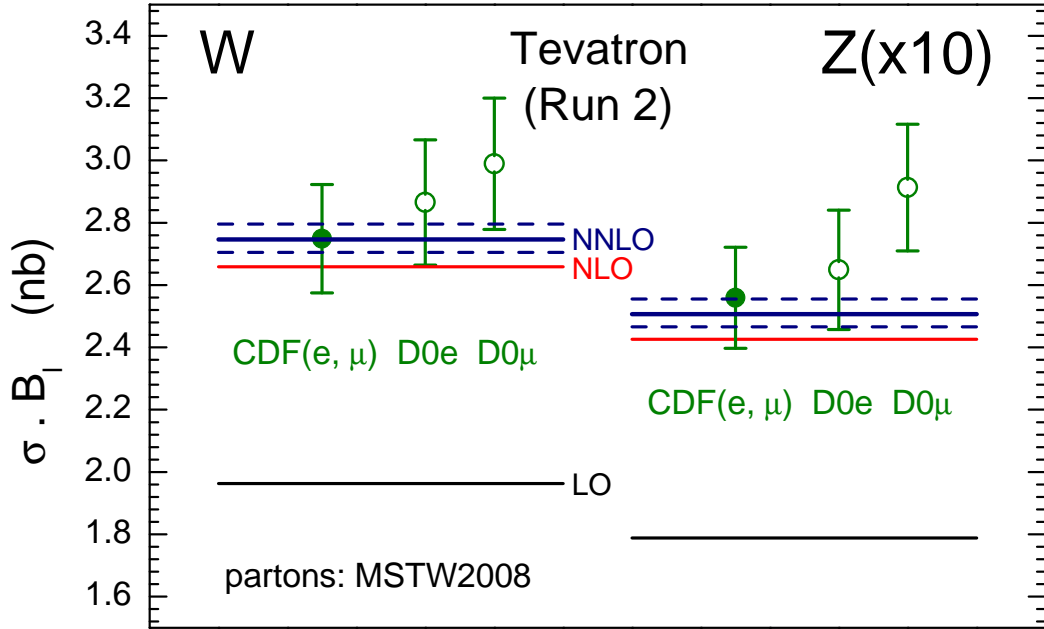


Figure 66: $W \equiv W^+ + W^-$ and Z total cross sections compared to Tevatron data [196, 197], where the dashed lines show the one-sigma PDF uncertainties on the NNLO predictions. The luminosity uncertainty of $\sim \pm 6\%$ is not included on the experimental data points.




Figure 67: $W \equiv W^+ + W^-$ and Z total cross sections at the Tevatron ($\sqrt{s} = 1.96$ TeV). The error ellipses, calculated using the one-sigma error sets, are shown for the MSTW 2008 NLO and NNLO PDFs. We show the CDF Run II measurement of the ratio $R_{WZ} = 10.84 \pm 0.15(\text{stat.}) \pm 0.14(\text{syst.})$ [196] with statistical and systematic uncertainties added in quadrature to get the overall uncertainty band.




Figure 68: $W \equiv W^+ + W^-$ and Z total cross sections at the LHC ($\sqrt{s} = 14$ TeV). The error ellipses, calculated using the one-sigma error sets, are shown for the MSTW 2008 NLO and NNLO PDFs.

Ratio to MSTW 2008	Tevatron		LHC	
	σ_W	σ_Z	σ_W	σ_Z
MRST 2006 NNLO	1.00	1.01	0.99	1.00
MRST 2004 NNLO	0.99	1.00	0.93	0.94
Alekhin02 NNLO	1.02	1.04	0.97	0.96
MRST 2006 NLO (unpublished)	0.99	1.00	1.00	1.01
MRST 2004 NLO	0.99	1.00	0.97	0.98
Alekhin02 NLO	1.03	1.05	1.01	1.00
CTEQ6.6 NLO	0.98	0.99	1.02	1.02

Table 10: Ratios of predictions for $W \equiv W^+ + W^-$ and Z total cross sections at the Tevatron ($\sqrt{s} = 1.96$ TeV) and LHC ($\sqrt{s} = 14$ TeV) calculated using the central values of different PDF sets, with respect to those from the MSTW 2008 sets.

a variety of other PDF sets. For the LHC predictions in Fig. 68, we show lines of constant $R_{WZ} = B_{\ell\nu} \cdot \sigma_W / B_{\ell\bar{\ell}} \cdot \sigma_Z$, and on the corresponding Tevatron plot in Fig. 67 we overlay a recent measurement of $R_{WZ} = 10.84 \pm 0.15(\text{stat.}) \pm 0.14(\text{syst.})$ from CDF [196].

The ratios of predictions for W and Z cross sections calculated using the central values of different PDF sets, with respect to those from the MSTW 2008 sets, are given in Table 10. The difference between the predictions of the various PDF sets can be understood by considering the values of the corresponding PDFs in the relevant x and Q^2 region. We have already noted that the relevant PDFs in the MRST 2006 and MSTW 2008 sets are very similar, see Fig. 61, and this is reflected in the similarity of the W and Z cross sections. We have also noted that a more careful treatment of the heavy quark (c and b) distributions at NNLO in the MRST 2006 and MSTW 2008 sets gave rise to a significant increase in the light quark distributions at high Q^2 , and this explains the increase in the W and Z cross sections compared to MRST 2004.³² The CTEQ6.6 NLO predictions are close to those of MSTW 2008 NLO at the Tevatron, but some 2% higher at the LHC. This can be traced back to the more rapid evolution of the CTEQ quarks at small x , driven by a (larger) gluon that, unlike MSTW 2008, is constrained to be positive at Q_0^2 , see Fig. 62. The Alekhin NLO and NNLO predictions for the ratio R_{WZ} at the LHC are larger than the predictions from all other PDF sets, due to the small strange quark distribution in the Alekhin PDF sets; see Section. 14.

We have so far considered total $W = W^+ + W^-$ production, but of course at the LHC $\sigma_{W^+} \neq \sigma_{W^-}$, and the ratio $R_{\pm} \equiv \sigma_{W^+}/\sigma_{W^-}$ is also of interest from the PDF perspective. Assuming a diagonal CKM matrix we have, at leading order,

$$R_{\pm} \approx \frac{u(x_1)\bar{d}(x_2) + c(x_1)\bar{s}(x_2) + (1 \leftrightarrow 2)}{d(x_1)\bar{u}(x_2) + s(x_1)\bar{c}(x_2) + (1 \leftrightarrow 2)}. \quad (99)$$

For central production, which dominates the total cross sections, we have $x_1 \sim x_2 \sim M_W/\sqrt{s}$ and $Q \sim M_W$. In the limit $\sqrt{s} \rightarrow \infty$, $R_{\pm} \rightarrow 1$ since the quark and antiquark distributions become equal in this limit. The fact that $R_{\pm} \neq 1$ at LHC is a reflection of the residual differences,

³²A similar effect was observed at NLO in the evolution of CTEQ6.1 to CTEQ6.5 and CTEQ6.6 [84, 87].

LHC, $\sqrt{s} = 10$ TeV	$B_{\ell\nu} \cdot \sigma_{W^+}$ (nb)	$B_{\ell\nu} \cdot \sigma_{W^-}$ (nb)	R_{\pm}
MSTW 2008 LO	$7.35^{+0.08}_{-0.12}$ (+1.1%) (-1.6%)	$5.22^{+0.06}_{-0.09}$ (+1.1%) (-1.7%)	$1.408^{+0.015}_{-0.012}$ (+1.0%) (-0.8%)
MSTW 2008 NLO	$8.62^{+0.18}_{-0.14}$ (+2.1%) (-1.7%)	$6.30^{+0.14}_{-0.11}$ (+2.2%) (-1.7%)	$1.367^{+0.012}_{-0.010}$ (+0.9%) (-0.7%)
MSTW 2008 NNLO	$8.88^{+0.15}_{-0.15}$ (+1.7%) (-1.6%)	$6.47^{+0.11}_{-0.11}$ (+1.7%) (-1.6%)	$1.373^{+0.012}_{-0.010}$ (+0.8%) (-0.7%)

LHC, $\sqrt{s} = 14$ TeV	$B_{\ell\nu} \cdot \sigma_{W^+}$ (nb)	$B_{\ell\nu} \cdot \sigma_{W^-}$ (nb)	R_{\pm}
MSTW 2008 LO	$10.69^{+0.14}_{-0.19}$ (+1.3%) (-1.8%)	$7.83^{+0.10}_{-0.14}$ (+1.2%) (-1.8%)	$1.366^{+0.013}_{-0.010}$ (+0.9%) (-0.8%)
MSTW 2008 NLO	$12.06^{+0.24}_{-0.21}$ (+2.0%) (-1.8%)	$9.11^{+0.19}_{-0.16}$ (+2.1%) (-1.6%)	$1.325^{+0.011}_{-0.009}$ (+0.8%) (-0.7%)
MSTW 2008 NNLO	$12.39^{+0.22}_{-0.21}$ (+1.8%) (-1.7%)	$9.33^{+0.16}_{-0.16}$ (+1.7%) (-1.7%)	$1.328^{+0.011}_{-0.009}$ (+0.8%) (-0.7%)

Table 11: Predictions for W^+ and W^- total cross sections at the LHC, including the one-sigma PDF uncertainties, and their ratio R_{\pm} . We take $\mu_R = \mu_F = M_W$.

particularly between u and \bar{u} and between d and \bar{d} , due to the valence quarks. Fig. 69 shows various predictions for the W^+ and W^- total cross sections at the LHC. Lines of constant R_{\pm} are also superimposed. The MSTW 2008 predictions are listed in Table 11.

Comparing Fig. 69 with the corresponding Fig. 68 for W and Z production, we see that there is less correlation between W^+ and W^- than between W and Z . This is because the combinations of u and d quark and antiquark distributions probed in $W \equiv W^+ + W^-$ and Z production are more similar than in W^+ and W^- separately. It is also interesting that the prediction for R_{\pm} has decreased significantly in going from MRST 2006 to MSTW 2008, due to a change in the u/d ratio resulting from the addition of the new Tevatron W and Z data and neutrino DIS data, as discussed in Section 11.

Note that unlike R_{WZ} , which additionally depends on electroweak parameter and branching ratio values, the overwhelmingly dominant uncertainties in the theoretical predictions for R_{\pm} are those due to the PDFs. The experimental measurement should also be very precise, since it is simply a matter of comparing the number of ℓ^+ and ℓ^- events in a sample of $W \rightarrow \ell\nu$ events. From Fig. 69, we see that a measurement of R_{\pm} with an error of less than 1% at the LHC will further constrain the parton distributions, particularly the u/d ratio.

In Fig. 70 we show the ratio of the parton luminosities

$$\frac{\partial \mathcal{L}_{ab}}{\partial M_X^2} = \frac{1}{s} \int_{\tau}^1 \frac{dx}{x} f_a(x, M_X^2) f_b(\tau/x, M_X^2), \quad \tau = \frac{M_X^2}{s} \quad (100)$$

at $\sqrt{s} = 10$ TeV compared to $\sqrt{s} = 14$ TeV at the LHC for $ab = gg, \sum_{q=u,\dots,b} q\bar{q}$, using the MSTW 2008 NLO PDFs including the one-sigma uncertainty bands. The PDF uncertainty on the ratio of the W total cross section at 10 TeV compared to 14 TeV is of order $\pm 0.3\%$, and similarly for the ratio of Z total cross sections. The PDF uncertainty is so small because it is only sensitive to the *slope* (rather than absolute values) of the PDFs in the region between $x = 0.006$ and $x = 0.01$ (at central rapidity). Other theoretical uncertainties, such as the choice of electroweak parameters, should also cancel out in this ratio. Experimentally, the accuracy of the measurement of this ratio will depend on whether the *relative* machine luminosity at 10 TeV and 14 TeV can be measured with high precision.




Figure 69: W^+ and W^- total cross sections at the LHC ($\sqrt{s} = 14$ TeV). The error ellipses, calculated using the one-sigma error sets, are shown for the MSTW 2008 NLO and NNLO PDFs.




Figure 70: Ratio of parton luminosities at $\sqrt{s} = 10$ TeV compared to $\sqrt{s} = 14$ TeV at the LHC, using the MSTW 2008 NLO PDFs including the one-sigma uncertainty bands.

16 Conclusions

We have presented the new MSTW 2008 parton distribution functions (PDFs) at leading-order (LO), next-to-leading order (NLO) and next-to-next-leading order (NNLO). In each case we include eigenvector sets spanning the uncertainty of the PDFs defined at a fixed value of $\alpha_S(M_Z^2)$. These sets are a major update to all our previously available sets, and incorporate the maximum amount of information from deep-inelastic scattering (DIS) and other hard-scattering data. We have also introduced a number of new theoretical refinements to produce what are arguably the most complete and reliable sets available. The PDFs are specifically designed for use in precision cross section predictions and uncertainties at the LHC. The grids and interpolation code (in Fortran, C++ and Mathematica formats) are available from HEPFORGE [27], and are also made available as part of the LHAPDF package [198] from version 5.7.0.

The principal new data sets included in the analysis are:

- neutrino structure function (F_2 and xF_3) data from NuTeV [37] and CHORUS [38];
- neutrino dimuon cross sections from CCFR and NuTeV [39], which constrain the strange quarks and antiquarks, allowing these to be fit directly for the first time, instead of using the previous input assumption that $s = \bar{s} = \kappa(\bar{u} + \bar{d})/2$ with $\kappa \approx 0.5$;
- the Z rapidity distribution [52, 53] and the lepton charge asymmetry from W decays [48, 49] measured at the Tevatron Run II, both of which constrain mainly the down quark distribution;
- the inclusive jet production data from the Tevatron Run II [54, 56], which prefer a smaller high- x gluon³³ than the previous Run I data;
- updated HERA data on $F_2^{c\bar{c}}$ [41–47] and on inclusive jet production in DIS [57–59].

In addition to including new data, there are a number of important developments in the fitting procedure.

- Our treatment of heavy flavours, carried out in a general-mass variable flavour number scheme (GM-VFNS), is the most complete available, particularly at NNLO where the improved treatment, first implemented in the MRST 2006 analysis [21], has a significant quantitative effect. At NLO we estimate the possible theoretical uncertainty on the parton distributions due to variations in the specific choice of GM-VFNS to be around 2%, though this should decrease at NNLO.
- We improve and extend our input parameterisation, not only introducing two new free parameters for each of $s + \bar{s}$ and $s - \bar{s}$, but using one more in the determination of the uncertainty on the gluon distribution, and making a better choice for the down valence quark distribution.

³³The smaller gluon distribution at high x has important consequences. For example, the predicted Tevatron cross section for production via gluon–gluon fusion of a Standard Model Higgs boson with mass $M_H = 170$ GeV decreases by almost 15% compared to the result using the previous MRST 2006 NNLO PDFs [199, 200] (for implications, see Ref. [201]), although this is also partly due to the smaller value of α_S for MSTW 2008 NNLO.

- We have designed and implemented a new treatment of PDF uncertainties, doing away with the conventional notion of an overall fixed tolerance $\Delta\chi^2_{\text{global}} = 50$ (for a 90% confidence-level limit) and replacing it with a new “dynamical tolerance” in which the tolerance for each independent eigenvector direction is determined from the condition that all data sets should be described within their 90% (or 68%) confidence-level limits. In the same context, we have also included uncertainties from data set normalisations for the first time. This can have an important effect on the PDF uncertainties, particularly for light quarks in the range of x and Q^2 where they are needed for collider phenomenology. It also means that, despite the presence of more constraining data in the fit, the PDF uncertainties have got slightly bigger in some cases.
- We have replaced the unconventional choice of α_S used in the MRST fits, which is not strictly correct at NNLO, with the exact solution of the renormalisation group equation (5), as used in recent public DGLAP evolution codes. This has enabled our evolution code to be checked against the PEGASUS [68] and HOPPET [69] codes for the first time.
- We use the FASTNLO package [75] for NLO (and for Tevatron jets also for approximate NNLO) calculations of jet cross sections directly in the fit. We have also improved our calculation of K -factors for low-mass Drell–Yan lepton pair and W and Z boson production using the DYRAP/VRAP [167] and FEWZ [73] codes.

We have fitted a wide variety of data (~ 2700 data points in all) and overall the quality of the NLO and NNLO fits is similar and is perfectly acceptable, with $\chi^2/N_{\text{pts.}} \sim 1$ for almost all data sets fitted. As has been noticed in previous studies, however, the quality of the LO fit is poor, and there are a number of examples where the data really do require significant NLO corrections. We therefore do not recommend that the LO set be used, for example, in leading-order parton-shower Monte Carlo event generators. A modified LO set (see Refs. [66, 67]) suitable for this purpose will be produced in a subsequent study.

Each of our best-fit PDF sets comes with a corresponding set of 40 error PDFs (compared to 30 in previous MRST versions), which can be used to estimate the PDF uncertainty on any physical quantity. It is important to remember that these PDF uncertainties come from the experimental errors on input data only. They do not include additional theory errors, coming, for example, from missing higher-order terms in the perturbation series. This issue will be addressed in a future study. While the convergence of the series is an important and outstanding issue for some quantities studied, for example, small x structure functions, it does not appear to be a problem in other cases, for example, $\sigma_{W,Z}$ at the hadron colliders.

In subsequent studies we will return to some of the more theoretical issues, including the dependence of the fit on the value of α_S and the heavy quark (c, b) masses, and the impact on the fits of including $\mathcal{O}(\alpha)$ QED corrections, as was done in Ref. [19]. We also intend to investigate W , Z and Drell–Yan lepton pair production at the LHC in more detail, with a particular focus on extracting information on the parton distributions at very small x .

In the meantime, these distributions are obtained from the most up-to-date and complete global fit yet performed, and we believe they are the most reliable sets available as we enter the LHC era.

Acknowledgements

We thank C. Gwenlan, K. Hatakeyama, T. Kluge, M. Lancaster, D. Mason, P. Nadolsky, F. Petriello, M. Wobisch and J. Zhu for valuable information, and S. Alekhin, A. Cooper-Sarkar, S. Forte, J. Huston, W.-K. Tung and A. Vogt for useful discussions. R.S.T. acknowledges the Royal Society for the award of a University Research Fellowship. G.W. acknowledges the UK Science and Technology Facilities Council for the award of a Responsive Research Associate position.

References

- [1] A. D. Martin, R. G. Roberts and W. J. Stirling, Phys. Rev. D **37** (1988) 1161.
- [2] A. D. Martin, R. G. Roberts and W. J. Stirling, Phys. Lett. B **206** (1988) 327.
- [3] P. N. Harriman, A. D. Martin, W. J. Stirling and R. G. Roberts, Phys. Rev. D **42** (1990) 798.
- [4] J. Kwiecinski, A. D. Martin, W. J. Stirling and R. G. Roberts, Phys. Rev. D **42** (1990) 3645.
- [5] A. D. Martin, W. J. Stirling and R. G. Roberts, Phys. Rev. D **47** (1993) 867.
- [6] A. D. Martin, W. J. Stirling and R. G. Roberts, Phys. Lett. B **306** (1993) 145 [Erratum-
ibid. B **309** (1993) 492].
- [7] A. D. Martin, W. J. Stirling and R. G. Roberts, Phys. Rev. D **50** (1994) 6734 [arXiv:hep-
ph/9406315].
- [8] A. D. Martin, W. J. Stirling and R. G. Roberts, Phys. Lett. B **354** (1995) 155 [arXiv:hep-
ph/9502336].
- [9] E. W. N. Glover, A. D. Martin, R. G. Roberts and W. J. Stirling, Phys. Lett. B **381**
(1996) 353 [arXiv:hep-ph/9603327].
- [10] A. D. Martin, R. G. Roberts and W. J. Stirling, Phys. Lett. B **387** (1996) 419 [arXiv:hep-
ph/9606345].
- [11] A. D. Martin, R. G. Roberts, M. G. Ryskin and W. J. Stirling, Eur. Phys. J. C **2** (1998)
287 [arXiv:hep-ph/9612449].
- [12] A. D. Martin, R. G. Roberts, W. J. Stirling and R. S. Thorne, Eur. Phys. J. C **4** (1998)
463 [arXiv:hep-ph/9803445].
- [13] A. D. Martin, R. G. Roberts, W. J. Stirling and R. S. Thorne, Eur. Phys. J. C **14** (2000)
133 [arXiv:hep-ph/9907231].
- [14] A. D. Martin, R. G. Roberts, W. J. Stirling and R. S. Thorne, Eur. Phys. J. C **23** (2002)
73 [arXiv:hep-ph/0110215].
- [15] A. D. Martin, R. G. Roberts, W. J. Stirling and R. S. Thorne, Phys. Lett. B **531** (2002)
216 [arXiv:hep-ph/0201127].
- [16] A. D. Martin, R. G. Roberts, W. J. Stirling and R. S. Thorne, Eur. Phys. J. C **28** (2003)
455 [arXiv:hep-ph/0211080].

- [17] A. D. Martin, R. G. Roberts, W. J. Stirling and R. S. Thorne, Eur. Phys. J. C **35** (2004) 325 [arXiv:hep-ph/0308087].
- [18] A. D. Martin, R. G. Roberts, W. J. Stirling and R. S. Thorne, Phys. Lett. B **604**, 61 (2004) [arXiv:hep-ph/0410230].
- [19] A. D. Martin, R. G. Roberts, W. J. Stirling and R. S. Thorne, Eur. Phys. J. C **39** (2005) 155 [arXiv:hep-ph/0411040].
- [20] A. D. Martin, W. J. Stirling and R. S. Thorne, Phys. Lett. B **636** (2006) 259 [arXiv:hep-ph/0603143].
- [21] A. D. Martin, W. J. Stirling, R. S. Thorne and G. Watt, Phys. Lett. B **652** (2007) 292 [arXiv:0706.0459 [hep-ph]].
- [22] W. L. van Neerven and A. Vogt, Nucl. Phys. B **568** (2000) 263 [arXiv:hep-ph/9907472].
- [23] W. L. van Neerven and A. Vogt, Nucl. Phys. B **588** (2000) 345 [arXiv:hep-ph/0006154].
- [24] W. L. van Neerven and A. Vogt, Phys. Lett. B **490** (2000) 111 [arXiv:hep-ph/0007362].
- [25] S. Moch, J. A. M. Vermaseren and A. Vogt, Nucl. Phys. B **688** (2004) 101 [arXiv:hep-ph/0403192].
- [26] A. Vogt, S. Moch and J. A. M. Vermaseren, Nucl. Phys. B **691** (2004) 129 [arXiv:hep-ph/0404111].
- [27] <http://projects.hepforge.org/mstwpdf/>
- [28] A. D. Martin, W. J. Stirling and R. S. Thorne, Phys. Lett. B **635** (2006) 305 [arXiv:hep-ph/0601247].
- [29] C. Adloff *et al.* [H1 Collaboration], Phys. Lett. B **393** (1997) 452 [arXiv:hep-ex/9611017].
- [30] N. Gogitidze, J. Phys. G **28** (2002) 751 [arXiv:hep-ph/0201047].
- [31] E. M. Lobodzinska [H1 Collaboration], arXiv:hep-ph/0311180.
- [32] A. C. Benvenuti *et al.* [BCDMS Collaboration], Phys. Lett. B **223** (1989) 485.
- [33] M. Arneodo *et al.* [New Muon Collaboration], Nucl. Phys. B **483** (1997) 3 [arXiv:hep-ph/9610231].
- [34] L. W. Whitlow, S. Rock, A. Bodek, E. M. Riordan and S. Dasu, Phys. Lett. B **250** (1990) 193.
- [35] C. Adloff *et al.* [H1 Collaboration], Eur. Phys. J. C **30** (2003) 1 [arXiv:hep-ex/0304003].

- [36] S. Chekanov *et al.* [ZEUS Collaboration], Eur. Phys. J. C **32** (2003) 1 [arXiv:hep-ex/0307043].
- [37] M. Tzanov *et al.* [NuTeV Collaboration], Phys. Rev. D **74** (2006) 012008 [arXiv:hep-ex/0509010].
- [38] G. Onengut *et al.* [CHORUS Collaboration], Phys. Lett. B **632** (2006) 65.
- [39] M. Goncharov *et al.* [NuTeV Collaboration], Phys. Rev. D **64** (2001) 112006 [arXiv:hep-ex/0102049].
- [40] A. O. Bazarko *et al.* [CCFR Collaboration], Z. Phys. C **65**, 189 (1995) [arXiv:hep-ex/9406007].
- [41] C. Adloff *et al.* [H1 Collaboration], Z. Phys. C **72** (1996) 593 [arXiv:hep-ex/9607012].
- [42] C. Adloff *et al.* [H1 Collaboration], Phys. Lett. B **528** (2002) 199 [arXiv:hep-ex/0108039].
- [43] A. Aktas *et al.* [H1 Collaboration], Eur. Phys. J. C **45** (2006) 23 [arXiv:hep-ex/0507081].
- [44] A. Aktas *et al.* [H1 Collaboration], Eur. Phys. J. C **40** (2005) 349 [arXiv:hep-ex/0411046].
- [45] J. Breitweg *et al.* [ZEUS Collaboration], Eur. Phys. J. C **12** (2000) 35 [arXiv:hep-ex/9908012].
- [46] S. Chekanov *et al.* [ZEUS Collaboration], Phys. Rev. D **69** (2004) 012004 [arXiv:hep-ex/0308068].
- [47] S. Chekanov *et al.* [ZEUS Collaboration], JHEP **0707** (2007) 074 [arXiv:0704.3562 [hep-ex]].
- [48] D. Acosta *et al.* [CDF Collaboration], Phys. Rev. D **71** (2005) 051104 [arXiv:hep-ex/0501023].
- [49] V. M. Abazov *et al.* [DØ Collaboration], Phys. Rev. D **77** (2008) 011106 [arXiv:0709.4254 [hep-ex]].
- [50] V. M. Abazov *et al.* [D0 Collaboration], Phys. Rev. Lett. **101** (2008) 211801 [arXiv:0807.3367 [hep-ex]].
- [51] F. Abe *et al.* [CDF Collaboration], Phys. Rev. Lett. **81** (1998) 5754 [arXiv:hep-ex/9809001].
- [52] J. Han *et al.* [CDF Collaboration], “ $d\sigma/dy$ distribution of Drell–Yan dielectron pairs,” Public Note, May 2008, <http://www-cdf.fnal.gov/physics/ewk/2008/dszdy/>.
- [53] V. M. Abazov *et al.* [DØ Collaboration], Phys. Rev. D **76** (2007) 012003 [arXiv:hep-ex/0702025].

- [54] A. Abulencia *et al.* [CDF - Run II Collaboration], Phys. Rev. D **75** (2007) 092006 [Erratum-
ibid. D **75** (2007) 119901] [arXiv:hep-ex/0701051].
- [55] T. Aaltonen *et al.* [CDF Collaboration], Phys. Rev. D **78** (2008) 052006 [Erratum-ibid. D
79 (2009) 119902] [arXiv:0807.2204v4 [hep-ex]].
- [56] V. M. Abazov *et al.* [DØ Collaboration], Phys. Rev. Lett. **101** (2008) 062001
[arXiv:0802.2400 [hep-ex]].
- [57] S. Chekanov *et al.* [ZEUS Collaboration], Phys. Lett. B **547** (2002) 164 [arXiv:hep-
ex/0208037].
- [58] S. Chekanov *et al.* [ZEUS Collaboration], Nucl. Phys. B **765** (2007) 1 [arXiv:hep-
ex/0608048].
- [59] A. Aktas *et al.* [H1 Collaboration], Phys. Lett. B **653** (2007) 134 [arXiv:0706.3722 [hep-
ex]].
- [60] W. L. van Neerven and E. B. Zijlstra, Phys. Lett. B **272** (1991) 127.
- [61] E. B. Zijlstra and W. L. van Neerven, Phys. Lett. B **273** (1991) 476.
- [62] E. B. Zijlstra and W. L. van Neerven, Phys. Lett. B **297** (1992) 377.
- [63] E. B. Zijlstra and W. L. van Neerven, Nucl. Phys. B **383** (1992) 525.
- [64] S. Moch, J. A. M. Vermaseren and A. Vogt, Phys. Lett. B **606** (2005) 123 [arXiv:hep-
ph/0411112].
- [65] J. A. M. Vermaseren, A. Vogt and S. Moch, Nucl. Phys. B **724** (2005) 3 [arXiv:hep-
ph/0504242].
- [66] A. Sherstnev and R. S. Thorne, Eur. Phys. J. C **55** (2008) 553 [arXiv:0711.2473 [hep-ph]].
- [67] A. Sherstnev and R. S. Thorne, arXiv:0807.2132 [hep-ph].
- [68] A. Vogt, Comput. Phys. Commun. **170** (2005) 65 [arXiv:hep-ph/0408244].
- [69] G. P. Salam and J. Rojo, Comput. Phys. Commun. **180** (2009) 120 [arXiv:0804.3755
[hep-ph]].
- [70] W. Giele *et al.*, arXiv:hep-ph/0204316.
- [71] M. Dittmar *et al.*, arXiv:hep-ph/0511119.
- [72] R. S. Thorne, Phys. Rev. D **73** (2006) 054019 [arXiv:hep-ph/0601245].
- [73] K. Melnikov and F. Petriello, Phys. Rev. D **74** (2006) 114017 [arXiv:hep-ph/0609070];
FEWZ code from <http://www.phys.hawaii.edu/~kirill/FEHiP.htm>.

- [74] C. Balazs and C. P. Yuan, Phys. Rev. D **56** (1997) 5558 [arXiv:hep-ph/9704258];
RESBOS code from <http://hep.pa.msu.edu/resum/>.
- [75] T. Kluge, K. Rabbertz and M. Wobisch, arXiv:hep-ph/0609285;
FASTNLO code from <http://projects.hepforge.org/fastnlo/>.
- [76] Z. Nagy, Phys. Rev. Lett. **88** (2002) 122003 [arXiv:hep-ph/0110315].
- [77] Z. Nagy, Phys. Rev. D **68** (2003) 094002 [arXiv:hep-ph/0307268].
- [78] S. Catani, D. de Florian, G. Rodrigo and W. Vogelsang, Phys. Rev. Lett. **93** (2004) 152003
[arXiv:hep-ph/0404240].
- [79] J. H. Kuhn, M. Steinhauser and C. Sturm, Nucl. Phys. B **778** (2007) 192 [arXiv:hep-ph/0702103].
- [80] J. H. Kuhn, M. Steinhauser and C. Sturm, arXiv:0705.2335 [hep-ph].
- [81] E. Laenen, S. Riemersma, J. Smith and W. L. van Neerven, Nucl. Phys. B **392**, 162 (1993).
- [82] B. W. Harris and J. Smith, Nucl. Phys. B **452** (1995) 109 [arXiv:hep-ph/9503484].
- [83] M. Buza, Y. Matiounine, J. Smith and W. L. van Neerven, Eur. Phys. J. C **1** (1998) 301
[arXiv:hep-ph/9612398].
- [84] W. K. Tung, H. L. Lai, A. Belyaev, J. Pumplin, D. Stump and C. P. Yuan, JHEP **0702**
(2007) 053 [arXiv:hep-ph/0611254].
- [85] M. A. G. Aivazis, J. C. Collins, F. I. Olness and W. K. Tung, Phys. Rev. D **50** (1994)
3102 [arXiv:hep-ph/9312319].
- [86] R. S. Thorne and R. G. Roberts, Phys. Rev. D **57** (1998) 6871 [arXiv:hep-ph/9709442].
- [87] P. M. Nadolsky *et al.*, Phys. Rev. D **78**, 013004 (2008) [arXiv:0802.0007 [hep-ph]].
- [88] R. S. Thorne and W. K. Tung, arXiv:0809.0714 [hep-ph].
- [89] A. Chuvakin, J. Smith and W. L. van Neerven, Phys. Rev. D **61** (2000) 096004 [arXiv:hep-ph/9910250].
- [90] W. K. Tung, S. Kretzer and C. Schmidt, J. Phys. G **28** (2002) 983 [arXiv:hep-ph/0110247].
- [91] S. Kretzer, H. L. Lai, F. I. Olness and W. K. Tung, Phys. Rev. D **69** (2004) 114005
[arXiv:hep-ph/0307022].
- [92] E. Laenen and S. O. Moch, Phys. Rev. D **59** (1999) 034027 [arXiv:hep-ph/9809550].
- [93] S. Catani, M. Ciafaloni and F. Hautmann, Nucl. Phys. B **366** (1991) 135.

- [94] I. Bierenbaum, J. Blumlein and S. Klein, Acta Phys. Polon. B **39** (2008) 1531 [arXiv:0806.0451 [hep-ph]].
- [95] S. Alekhin and S. Moch, Phys. Lett. B **672**, 166 (2009) [arXiv:0811.1412 [hep-ph]].
- [96] M. Kramer, F. I. Olness and D. E. Soper, Phys. Rev. D **62** (2000) 096007 [arXiv:hep-ph/0003035].
- [97] M. Buza and W. L. van Neerven, Nucl. Phys. B **500** (1997) 301 [arXiv:hep-ph/9702242].
- [98] S. J. Brodsky, P. Hoyer, C. Peterson and N. Sakai, Phys. Lett. B **93** (1980) 451.
- [99] J. Pumplin, H. L. Lai and W. K. Tung, Phys. Rev. D **75** (2007) 054029 [arXiv:hep-ph/0701220].
- [100] R. S. Thorne, J. Phys. G **25** (1999) 1307 [arXiv:hep-ph/9902299].
- [101] J. C. Collins, Phys. Rev. D **58** (1998) 094002 [arXiv:hep-ph/9806259].
- [102] A. C. Benvenuti *et al.* [BCDMS Collaboration], Phys. Lett. B **237** (1990) 592.
- [103] M. Arneodo *et al.* [New Muon Collaboration], Nucl. Phys. B **487** (1997) 3 [arXiv:hep-ex/9611022].
- [104] M. R. Adams *et al.* [E665 Collaboration], Phys. Rev. D **54** (1996) 3006.
- [105] L. W. Whitlow, E. M. Riordan, S. Dasu, S. Rock and A. Bodek, Phys. Lett. B **282** (1992) 475.
- [106] L. W. Whitlow, Ph.D. thesis, Stanford University, 1990, SLAC-0357.
- [107] J. C. Webb, Ph.D. thesis, New Mexico State University, 2002, arXiv:hep-ex/0301031; Paul E. Reimer, private communication (for the radiative corrections).
- [108] R. S. Towell *et al.* [FNAL E866/NuSea Collaboration], Phys. Rev. D **64** (2001) 052002 [arXiv:hep-ex/0103030].
- [109] C. Adloff *et al.* [H1 Collaboration], Eur. Phys. J. C **21** (2001) 33 [arXiv:hep-ex/0012053].
- [110] C. Adloff *et al.* [H1 Collaboration], Eur. Phys. J. C **19** (2001) 269 [arXiv:hep-ex/0012052].
- [111] J. Breitweg *et al.* [ZEUS Collaboration], Eur. Phys. J. C **7** (1999) 609 [arXiv:hep-ex/9809005].
- [112] S. Chekanov *et al.* [ZEUS Collaboration], Eur. Phys. J. C **21** (2001) 443 [arXiv:hep-ex/0105090].
- [113] S. Chekanov *et al.* [ZEUS Collaboration], Eur. Phys. J. C **28** (2003) 175 [arXiv:hep-ex/0208040].

- [114] S. Chekanov *et al.* [ZEUS Collaboration], Phys. Rev. D **70** (2004) 052001 [arXiv:hep-ex/0401003].
- [115] I. Schienbein *et al.*, J. Phys. G **35** (2008) 053101 [arXiv:0709.1775 [hep-ph]].
- [116] Y. L. Dokshitzer, G. Marchesini and B. R. Webber, Nucl. Phys. B **469** (1996) 93 [arXiv:hep-ph/9512336].
- [117] M. Dasgupta and B. R. Webber, Phys. Lett. B **382** (1996) 273 [arXiv:hep-ph/9604388].
- [118] B. Badelek and J. Kwiecinski, Phys. Rev. D **50** (1994) 4 [arXiv:hep-ph/9401314].
- [119] A. d. R. Vargas Trevino [H1 Collaboration], “Measurement of the Inclusive ep Scattering Cross Section at low Q^2 and x at HERA”, presented at 15th International Workshop on Deep-Inelastic Scattering and Related Subjects (DIS2007), Munich, Germany, 16-20 April 2007, H1prelim-07-045. doi:10.3360/dis.2007.22
- [120] N. Kidonakis and J. F. Owens, Phys. Rev. D **63** (2001) 054019 [arXiv:hep-ph/0007268].
- [121] V. Blobel, proceedings of PHYSTAT2003: Statistical Problems in Particle Physics, Astrophysics, and Cosmology, Menlo Park, California, 8-11 September 2003.
- [122] A. Vogt, Nucl. Phys. Proc. Suppl. **79** (1999) 102 [arXiv:hep-ph/9906337].
- [123] R. Devenish and A. Cooper-Sarkar, “Deep inelastic scattering”, Oxford, UK: Univ. Pr. (2004).
- [124] H1 and ZEUS Collaborations, “Combination of H1 and ZEUS Deep Inelastic $e^\pm p$ Scattering Cross Section measurements”, submitted to The XXIII International Symposium on Lepton and Photon Interactions at High Energy (LP 2007), 13–18 August 2007, Daegu, Korea, H1prelim-07-007, ZEUS-prel-07-026.
- [125] W. K. Tung, talk in Structure Functions Working Group, XIV International Workshop on Deep Inelastic Scattering (DIS 2006), 20–24 April 2006, Tsukuba, Japan, <http://www-conf.kek.jp/dis06/program-general.htm>.
- [126] J. Pumplin, D. R. Stump, J. Huston, H. L. Lai, P. M. Nadolsky and W. K. Tung, JHEP **0207** (2002) 012 [arXiv:hep-ph/0201195].
- [127] F. James and M. Roos, Comput. Phys. Commun. **10** (1975) 343; F. James and M. Roos, MINUIT: Function Minimization and Error Analysis, CERN Program Library Long Writeup D506 (1994); F. James and M. Winkler, MINUIT User’s Guide: C++ Version (2004).
- [128] J. Pumplin, D. R. Stump and W. K. Tung, Phys. Rev. D **65** (2002) 014011 [arXiv:hep-ph/0008191].

- [129] K. Levenberg, “A Method for the Solution of Certain Non-Linear Problems in Least Squares”, *Quart. Appl. Math.* **2** (1944) 164.
- [130] D. W. Marquardt, “An Algorithm for Least-Squares Estimation of Nonlinear Parameters”, *SIAM J. Appl. Math.* **11** (1963) 431, doi:10.1137/0111030.
- [131] W. H. Press, B. P. Flannery, S. A. Teukolsky and W. T. Vetterling, *Numerical Recipes: The Art of Scientific Computing*, <http://www.nr.com/>.
- [132] C. Amsler *et al.* [Particle Data Group], *Phys. Lett. B* **667** (2008) 1.
- [133] D. Stump *et al.*, *Phys. Rev. D* **65** (2002) 014012 [arXiv:hep-ph/0101051].
- [134] J. Pumplin *et al.*, *Phys. Rev. D* **65** (2002) 014013 [arXiv:hep-ph/0101032].
- [135] W. T. Giele and S. Keller, *Phys. Rev. D* **58** (1998) 094023 [arXiv:hep-ph/9803393].
- [136] W. T. Giele, S. A. Keller and D. A. Kosower, arXiv:hep-ph/0104052.
- [137] R. D. Ball *et al.* [NNPDF Collaboration], *Nucl. Phys. B* **809** (2009) 1 [arXiv:0808.1231 [hep-ph]].
- [138] J. C. Collins and J. Pumplin, arXiv:hep-ph/0105207.
- [139] S. Alekhin, *Phys. Rev. D* **68** (2003) 014002 [arXiv:hep-ph/0211096].
- [140] S. Chekanov *et al.* [ZEUS Collaboration], *Eur. Phys. J. C* **42** (2005) 1 [arXiv:hep-ph/0503274].
- [141] M. Gluck, P. Jimenez-Delgado and E. Reya, *Eur. Phys. J. C* **53** (2008) 355 [arXiv:0709.0614 [hep-ph]].
- [142] A. Airapetian *et al.* [HERMES Collaboration], *Phys. Lett. B* **666** (2008) 446 [arXiv:0803.2993 [hep-ex]].
- [143] M. Dittmar *et al.*, arXiv:0901.2504 [hep-ph].
- [144] C. D. White and R. S. Thorne, *Phys. Rev. D* **75** (2007) 034005 [arXiv:hep-ph/0611204].
- [145] U. K. Yang *et al.* [CCFR/NuTeV Collaboration], *Phys. Rev. Lett.* **86** (2001) 2742 [arXiv:hep-ex/0009041].
- [146] M. Hirai, S. Kumano and T. H. Nagai, *Phys. Rev. C* **76**, 065207 (2007) [arXiv:0709.3038 [hep-ph]].
- [147] D. de Florian and R. Sassot, *Phys. Rev. D* **69** (2004) 074028 [arXiv:hep-ph/0311227].
- [148] M. Gluck, E. Reya and A. Vogt, *Eur. Phys. J. C* **5** (1998) 461 [arXiv:hep-ph/9806404].

- [149] K. J. Eskola, H. Paukkunen and C. A. Salgado, JHEP **0904** (2009) 065 [arXiv:0902.4154 [hep-ph]].
- [150] I. Schienbein, J. Y. Yu, C. Keppel, J. G. Morfin, F. Olness and J. F. Owens, Phys. Rev. D **77** (2008) 054013 [arXiv:0710.4897 [hep-ph]].
- [151] S. Alekhin, JETP Lett. **82** (2005) 628 [Pisma Zh. Eksp. Teor. Fiz. **82** (2005) 710] [arXiv:hep-ph/0508248].
- [152] J. Blumlein, H. Böttcher and A. Guffanti, Nucl. Phys. B **774** (2007) 182 [arXiv:hep-ph/0607200].
- [153] D. A. Mason, “Measurement of the strange–antistrange asymmetry at NLO in QCD from NuTeV dimuon data,” FERMILAB-THESIS-2006-01; “Final strange asymmetry results from NuTeV,” Proceedings of 14th International Workshop on Deep Inelastic Scattering (DIS 2006), Tsukuba, Japan, 20-24 Apr 2006; D. A. Mason *et al.*, Phys. Rev. Lett. **99** (2007) 192001.
- [154] S. Alekhin, S. Kulagin and R. Petti, Phys. Lett. B **675** (2009) 433 [arXiv:0812.4448 [hep-ph]].
- [155] F. Olness *et al.*, Eur. Phys. J. C **40** (2005) 145 [arXiv:hep-ph/0312323].
- [156] H. L. Lai, P. Nadolsky, J. Pumplin, D. Stump, W. K. Tung and C. P. Yuan, JHEP **0704** (2007) 089 [arXiv:hep-ph/0702268].
- [157] S. Kretzer, D. Mason and F. Olness, Phys. Rev. D **65** (2002) 074010 [arXiv:hep-ph/0112191].
- [158] P. D. B. Collins and T. P. Spiller, J. Phys. G **11** (1985) 1289.
- [159] J. Rojo *et al.* [NNPDF Collaboration], arXiv:0811.2288 [hep-ph].
- [160] T. Aaltonen *et al.* [CDF Collaboration], Phys. Rev. Lett. **100** (2008) 091803 [arXiv:0711.2901 [hep-ex]].
- [161] S. Davidson, S. Forte, P. Gambino, N. Rius and A. Strumia, JHEP **0202** (2002) 037 [arXiv:hep-ph/0112302].
- [162] G. P. Zeller *et al.* [NuTeV Collaboration], Phys. Rev. Lett. **88** (2002) 091802 [Erratum-ibid. **90** (2003) 239902] [arXiv:hep-ex/0110059].
- [163] G. P. Zeller *et al.* [NuTeV Collaboration], Phys. Rev. D **65** (2002) 111103 [Erratum-ibid. D **67** (2003) 119902] [arXiv:hep-ex/0203004].
- [164] K. S. McFarland and S. O. Moch, arXiv:hep-ph/0306052.
- [165] J. J. Aubert *et al.* [European Muon Collaboration], Nucl. Phys. B **213** (1983) 31.

- [166] E. Hoffmann and R. Moore, Z. Phys. C **20** (1983) 71.
- [167] C. Anastasiou, L. J. Dixon, K. Melnikov and F. Petriello, Phys. Rev. D **69** (2004) 094008 [arXiv:hep-ph/0312266];
VRAP code from <http://www.slac.stanford.edu/~lance/Vrap/>.
- [168] S. Alekhin, K. Melnikov and F. Petriello, Phys. Rev. D **74** (2006) 054033 [arXiv:hep-ph/0606237].
- [169] N. E. Adam, V. Halyo and S. A. Yost, JHEP **0805** (2008) 062 [arXiv:0802.3251 [hep-ph]].
- [170] N. E. Adam, V. Halyo, S. A. Yost and W. Zhu, JHEP **0809** (2008) 133 [arXiv:0808.0758 [hep-ph]].
- [171] T. Carli, G. P. Salam and F. Siegert, arXiv:hep-ph/0510324; T. Carli, D. Clements *et al.*, in preparation.
- [172] B. Y. Han *et al.* [CDF Collaboration], “Direct Measurement of W Boson Charge Asymmetry with 1 fb^{-1} of Run II data,” CDF note 8942; B. Y. Han, “Measurement of the W Boson Production Charge Asymmetry in $p\bar{p}$ Collisions,” FERMILAB-THESIS-2008-15.
- [173] T. Aaltonen *et al.* [The CDF Collaboration], arXiv:0901.2169 [hep-ex].
- [174] F. Abe *et al.* [CDF Collaboration], Phys. Rev. Lett. **77** (1996) 438 [arXiv:hep-ex/9601008].
- [175] H. L. Lai *et al.*, Phys. Rev. D **55** (1997) 1280 [arXiv:hep-ph/9606399].
- [176] B. Abbott *et al.* [DØ Collaboration], Phys. Rev. Lett. **86** (2001) 1707 [arXiv:hep-ex/0011036].
- [177] D. Stump, J. Huston, J. Pumplin, W. K. Tung, H. L. Lai, S. Kuhlmann and J. F. Owens, JHEP **0310** (2003) 046 [arXiv:hep-ph/0303013].
- [178] M. L. Mangano, arXiv:0802.0026 [hep-ph].
- [179] A. A. Affolder *et al.* [CDF Collaboration], Phys. Rev. D **64** (2001) 032001 [Erratum-ibid. D **65** (2002) 039903] [arXiv:hep-ph/0102074].
- [180] J. M. Campbell and R. K. Ellis, Phys. Rev. D **62** (2000) 114012 [arXiv:hep-ph/0006304];
MCFM code from <http://mcfm.fnal.gov/>.
- [181] C. Adloff *et al.* [H1 Collaboration], Eur. Phys. J. C **19** (2001) 289 [arXiv:hep-ex/0010054].
- [182] A. Aktas *et al.* [H1 Collaboration], JHEP **0710** (2007) 042 [arXiv:0708.3217 [hep-ex]].
- [183] F. D. Aaron *et al.* [H1 Collaboration], Phys. Lett. B **665**, 139 (2008) [arXiv:0805.2809 [hep-ex]].
- [184] V. Chekelian, arXiv:0810.5112 [hep-ex].

- [185] G. Watt, A. D. Martin and M. G. Ryskin, Phys. Lett. B **627** (2005) 97 [arXiv:hep-ph/0508093].
- [186] H. Kowalski, L. Motyka and G. Watt, Phys. Rev. D **74** (2006) 074016 [arXiv:hep-ph/0606272].
- [187] E. Stein, M. Meyer-Hermann, L. Mankiewicz and A. Schafer, Phys. Lett. B **376** (1996) 177 [arXiv:hep-ph/9601356].
- [188] R. S. Thorne, arXiv:0808.1845 [hep-ph].
- [189] G. Altarelli, R. D. Ball and S. Forte, Nucl. Phys. B **799** (2008) 199 [arXiv:0802.0032 [hep-ph]].
- [190] M. Ciafaloni, D. Colferai, G. P. Salam and A. M. Stasto, JHEP **0708**, 046 (2007) [arXiv:0707.1453 [hep-ph]].
- [191] J. Bartels, K. J. Golec-Biernat and K. Peters, Eur. Phys. J. C **17** (2000) 121 [arXiv:hep-ph/0003042].
- [192] G. Watt and H. Kowalski, Phys. Rev. D **78** (2008) 014016 [arXiv:0712.2670 [hep-ph]].
- [193] R. Hamberg, W. L. van Neerven and T. Matsuura, Nucl. Phys. B **359** (1991) 343 [Erratum-ibid. B **644** (2002) 403];
original code from <http://www.lorentz.leidenuniv.nl/~neerven/>.
- [194] R. V. Harlander and W. B. Kilgore, Phys. Rev. Lett. **88** (2002) 201801 [arXiv:hep-ph/0201206].
- [195] M. Cacciari, S. Frixione, M. L. Mangano, P. Nason and G. Ridolfi, JHEP **0809** (2008) 127 [arXiv:0804.2800 [hep-ph]].
- [196] A. Abulencia *et al.* [CDF Collaboration], J. Phys. G **34** (2007) 2457 [arXiv:hep-ex/0508029].
- [197] DØ Collaboration: DØNote 4403-CONF, DØnote 4750, DØnote 4573.
- [198] M. R. Whalley, D. Bourilkov and R. C. Group, arXiv:hep-ph/0508110;
LHAPDF code from <http://projects.hepforge.org/lhapdf/>.
- [199] C. Anastasiou, R. Boughezal and F. Petriello, JHEP **0904** (2009) 003 [arXiv:0811.3458 [hep-ph]].
- [200] D. de Florian and M. Grazzini, Phys. Lett. B **674** (2009) 291 [arXiv:0901.2427 [hep-ph]].
- [201] P. Bechtle, O. Brein, S. Heinemeyer, G. Weiglein and K. E. Williams, arXiv:0811.4169 [hep-ph].