3.1 The state of the art

The data, N³LO, theory uncertainties and the precision the LHC programme needs.

1. How to read the whitepaper

The Snowmass 2021 whitepaper *Proton structure at the precision frontier* [1] is the field's own account of where the determination of parton distribution functions (PDFs) stood in 2022 and what the next decade requires. This chapter is a guided reading of it. Every statement about the literature below is taken from the whitepaper and carries its section number, written "Sect. V B". The relations it states in words are derived here, and those derivations are ours.

Three points about the document help in reading it.

- **It is written by sections.** Each has its own leading authors, and quantitative projections "reflect the assumptions and methodologies used by the leading authors" (Sect. I B). A projected uncertainty is that group's estimate, made with that group's method.
- **It writes N2LO where this guide writes NNLO**, next-to-next-to-leading order. N³LO is one order beyond.
- **It organises the subject into three domains** (its Fig. 1): experiment, theory and statistics.

DOMAIN	WHITEPAPER	WHAT IS AT STAKE	HERE
Applications	Sect. II	Higgs physics, searches, effective-theory fits	sections 2 to 5
Experiment	Sect. III	LHC, Electron-Ion Collider, LHeC, forward experiments	section 6
Theory	Sect. IV	N³LO, electroweak corrections, resummation	section 7
Statistics and tools	Sect. V, IX, X	systematic and theory uncertainties and delivery (V), computing tools (IX), combination (X)	sections 4, 8, 9
Other targets	Sect. VI to VIII	lattice input, nuclei and mesons, transverse momentum	section 10, chapter 3.3

Its Table I measures progress since the Snowmass study of 2013. Then, NNLO was available for evolution, deep inelastic scattering (DIS) and vector-boson production. In 2021 it is available for all key processes, and N³LO for some.

Chapters 1.6 and 2.1 to 2.5 covered which process constrains which distribution and how a fit is built. This chapter adds five derived results: the uncertainty of a parton luminosity, the propagation of error sets, the effect of new data, the effect of a theory covariance matrix, and the uncertainty of a combination.

2. Why the distributions limit precision

This section identifies four measurements in which the PDF uncertainty is a leading term.

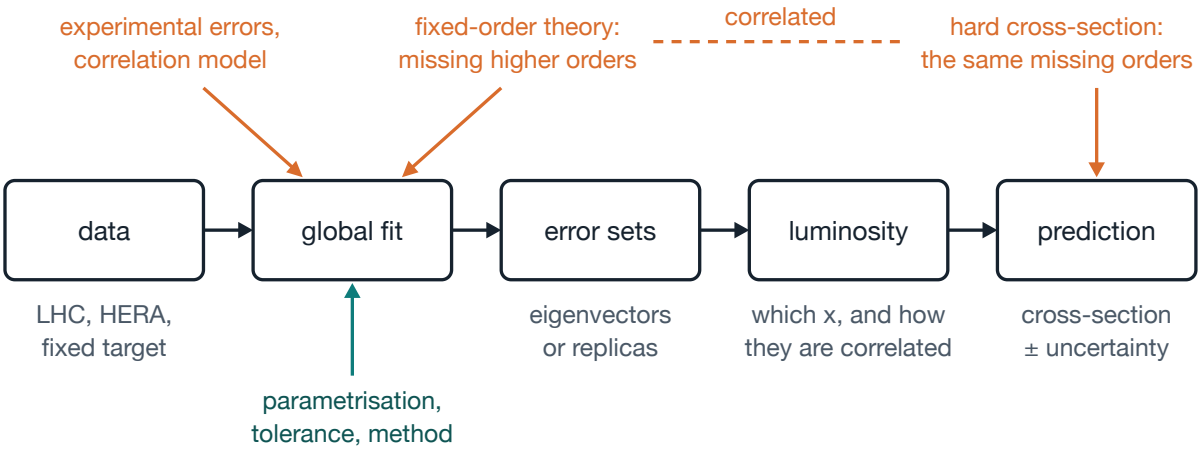

Figure 1. How an uncertainty reaches a prediction. Three inputs shape the fit: the experimental errors with their assumed correlations (section 6), the truncated perturbative theory (sections 7 and 8), and the choices of method (section 9). The error sets carry the result to a luminosity (section 3) and to a cross-section (section 4). The dashed line is the point of section 8: the theory in the fit and the theory in the prediction leave out related terms.

Higgs production

In the Standard Model the Higgs mass fixes every Higgs coupling, so a measured deviation would signal new physics. The test is only as sharp as the predicted cross-section (Sect. II B 1). For gluon fusion the partonic cross-section is known at N³LO. The whitepaper's Fig. 7 splits the remaining theory uncertainty into scale variation, electroweak corrections, quark-mass effects and two terms from the distributions. The first, $\delta(\text{PDF} + \alpha_S)$, is the propagated uncertainty of the distributions and of the coupling, combined in quadrature. The second exists because the cross-section is N³LO and the distributions are NNLO:

$$\delta(\text{PDF-TH}) = \frac{1}{2}\,\frac{\left|\sigma^{\text{NNLO}}_{\text{NNLO PDF}} - \sigma^{\text{NNLO}}_{\text{NLO PDF}}\right|}{\sigma^{\text{NNLO}}_{\text{NNLO PDF}}}. \tag{3.1.1}$$

The superscript is the order of the partonic cross-section and the subscript the order of the distributions. The effect of changing the order of the distributions is measured one order down, where both sets exist, and halved on the expectation that the next step is smaller (Sect. IV B). The whitepaper calls the prescription ad hoc and very conservative. It still gives several per cent for gluon fusion, comparable to $\delta(\text{PDF} + \alpha_S)$.

Electroweak parameters and the strong coupling

In recent measurements of the W mass and the weak mixing angle, the PDF uncertainty is one of the largest (Sect. III A). Two difficulties are specific to this use.

- **The parameter and the distributions are correlated.** A PDF set fitted with a parameter held fixed is biased towards that value when used to extract the parameter, as Moore explains for α_S [2]. Ideally the two are fitted together. A family of sets at different fixed values is not enough (Sect. V B).
- **Sets are compared, not combined.** The uncertainty is usually estimated by repeating the measurement with different global sets, which ignores the correlation between sets that share most of their data (section 9).

The coupling and the gluon enter hadronic cross-sections together, so each limits the other. The whitepaper quotes the lattice average $\alpha_S(m_Z) = 0.1184(8)$ (Sect. VI A) and a projection of ± 0.00018 for the LHeC, which it describes as a factor of 6 below the present world average (Sect. III C).

Searches at high mass

LHC bounds disfavour new vector bosons lighter than about 4–5 TeV (Sect. II B 2). Heavier ones are produced from partons at large x, where uncertainties are large. A second problem is circular (Sect. II B 3). The same high-energy data are used to fit the distributions, assuming the Standard Model, and to bound new interactions, assuming Standard Model distributions. A real deviation could be absorbed into the fit. Chapter 3.2 treats this in full.

---

3. The parton luminosity and its uncertainty

This section defines the quantity in which the whitepaper compares PDF sets, shows which momentum fractions it samples, and derives its uncertainty.

Definition

For partons i and j in two protons colliding at centre-of-mass energy $\sqrt{s}$, the whitepaper's Eq. (1) is

$$\mathcal{L}_{ij}(m_X^2, \mu_F^2) = \frac{1}{s}\,\frac{1}{1+\delta_{ij}} \int_\tau^1 \frac{dx}{x}\Big[f_i(x,\mu_F^2)\, f_j\Big(\frac{\tau}{x},\mu_F^2\Big) + (i \leftrightarrow j)\Big], \qquad \tau = \frac{m_X^2}{s}. \tag{3.1.2}$$

Here m_X is the invariant mass of the parton pair, $f_i(x,\mu_F^2)$ is the distribution of parton i at factorisation scale $\mu_F = m_X$, and δ_{ij} prevents double counting of identical partons. The luminosity of chapter 1.6 is $s\,\mathcal{L}_{q\bar{q}}$ for one flavour. The whitepaper compares $\mathcal{L}_{gg}$ and three sums over flavours:

$$\mathcal{L}_{q\bar{q}} = \sum_i \mathcal{L}_{q_i\bar{q}_i}, \qquad \mathcal{L}_{qq} = \sum_i \big(\mathcal{L}_{q_iq_i} + \mathcal{L}_{\bar{q}_i\bar{q}_i}\big), \qquad \mathcal{L}_{gq} = \sum_i \big(\mathcal{L}_{gq_i} + \mathcal{L}_{g\bar{q}_i}\big). \tag{3.1.3}$$

What it is the luminosity of

Insert the definition into the factorisation theorem of chapter 1.6 for a total cross-section, with partonic cross-sections $\hat{\sigma}_{ij}$ that depend on the pair's invariant mass (derivation D1):

$$\sigma = \sum_{i\le j} \int_0^s dm_X^2\; \mathcal{L}_{ij}(m_X^2,\mu_F^2)\; \hat{\sigma}_{ij}(m_X^2). \tag{3.1.4}$$

The sum runs over unordered pairs. $\mathcal{L}_{ij}\,dm_X^2$ is the number of parton pairs of type ij, per pair of protons, with invariant mass squared in dm_X^2. For a narrow state produced at lowest order, $\hat{\sigma} = \sigma_0\, m_X^2\, \delta(\hat{s} - m_X^2)$ with $\hat{s}$ the squared mass of the pair, and

$$\sigma = \sigma_0\, m_X^2\, \mathcal{L}_{ij}(m_X^2, m_X^2). \tag{3.1.5}$$

A ratio of luminosities from two sets is the ratio of their lowest-order predictions.

Which momentum fractions

Change variable to the rapidity y of the pair, with $x_{1,2} = \sqrt{\tau}\,e^{\pm y}$ as in chapter 1.6:

$$s\,\mathcal{L}_{gg} = \int_{-y_{\max}}^{y_{\max}} dy\; g\big(\sqrt{\tau}e^{y}\big)\, g\big(\sqrt{\tau}e^{-y}\big), \qquad y_{\max} = \ln\frac{1}{\sqrt{\tau}}. \tag{3.1.6}$$

For a Higgs boson of mass 125 GeV at $\sqrt{s} = 14$ TeV, $\sqrt{\tau} = 0.0089$ and $y_{\max} = 4.72$. The two gluons have $x_1 x_2 = \tau$: one sits above 0.0089 and the other below, except at $y = 0$. Within $|y| < 2.5$ they range from 7.3×10^{-4} to 0.11.

The spread over y depends on the shape of the gluon. For $x\,g \propto x^{-\lambda}(1-x)^{b}$ the integrand of Eq. (3.1.6) is proportional to $\tau^{-1-\lambda}\big[(1-x_1)(1-x_2)\big]^{b}$. The small-$x$ power has dropped out, and the integrand is flat in rapidity until one gluon reaches large x.

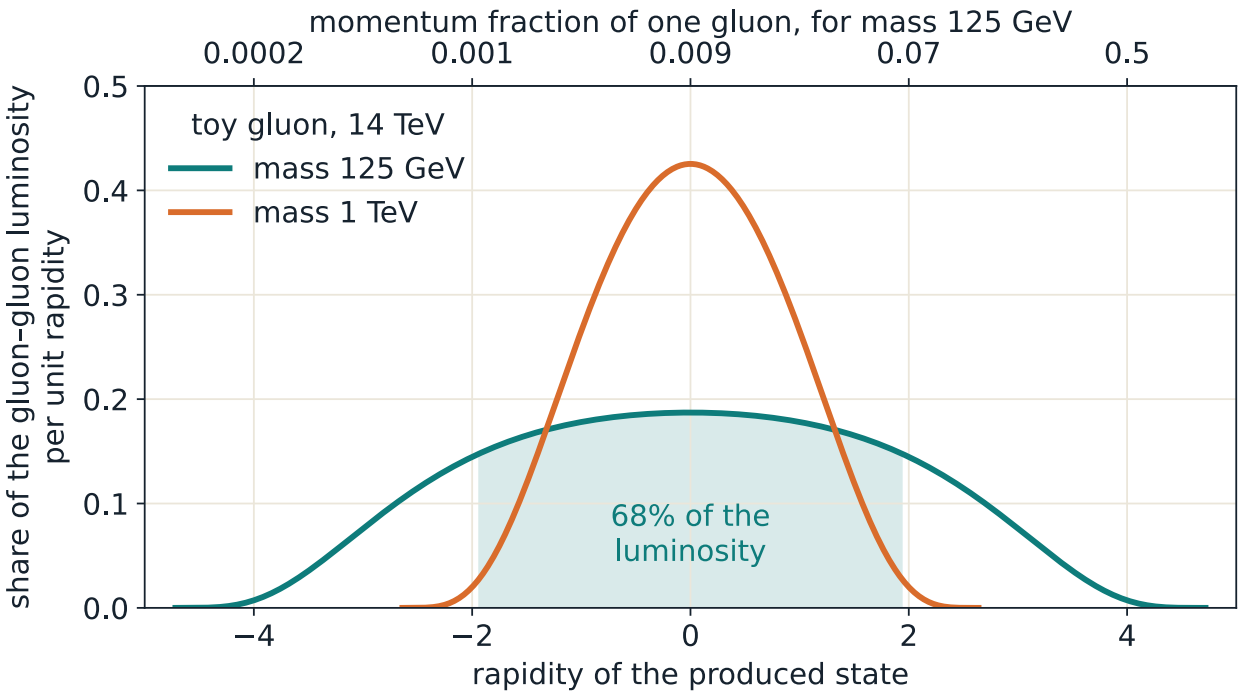

Figure 2. Where the gluon–gluon luminosity comes from, for a toy gluon $x\,g \propto x^{-0.3}(1-x)^{5}$ at $\sqrt{s} = 14$ TeV. The curves are the normalised integrand of Eq. (3.1.6). For $m_X = 125$ GeV, 68% of the luminosity lies within $|y| < 1.94$, where the gluons have $1.3 \times 10^{-3} < x < 0.062$ (upper axis). The shape is illustrative, not a fit.

For this toy shape, 37% of the luminosity at the Higgs mass comes from $|y| < 1$ and 83% from $|y| < 2.5$. At $m_X = 1$ TeV, $\sqrt{\tau} = 0.071$ and 74% comes from $|y| < 1$. At $m_X = 5$ TeV both partons have $x > \tau = 0.13$. A Higgs cross-section tests the gluon over nearly two decades of x, and a search at 5 TeV tests the large-x region, which the whitepaper calls the least known (Sect. IV D 1).

From a PDF uncertainty to a luminosity uncertainty

Let the gluon change by a small relative amount $r(x) = \delta g(x)/g(x)$. To first order (derivation D2),

$$\frac{\delta \mathcal{L}_{gg}}{\mathcal{L}_{gg}} = \int_{-y_{\max}}^{y_{\max}} dy\; w(y)\,\big[r(x_1) + r(x_2)\big], \qquad w(y) = \frac{g(x_1)\,g(x_2)}{s\,\mathcal{L}_{gg}}. \tag{3.1.7}$$

The weight $w(y)$ is the curve of Figure 2. In general $r(x_1)$ and $r(x_2)$ are two random quantities. Let each have standard deviation r and let their correlation coefficient be ρ. For a single pair of momentum fractions,

$$\frac{\delta \mathcal{L}_{gg}}{\mathcal{L}_{gg}} = r\,\sqrt{2\,(1+\rho)}. \tag{3.1.8}$$

> **A 1% gluon is not a 1% cross-section.** With a gluon uncertainty of 1%, Eq. (3.1.8) gives a luminosity uncertainty of 2.0% if the gluon moves coherently ($\rho = 1$), 1.4% if the two momentum fractions fluctuate independently ($\rho = 0$), and zero if they are fully anticorrelated ($\rho = -1$). A plot of a PDF with its band does not contain the correlations. The error sets do.

By Eq. (3.1.5) the same percentages apply to the lowest-order cross-section for $gg \to H$. That cross-section is also proportional to α_S^2, so the lattice uncertainty of section 2 contributes $2 \times 0.0008/0.1184 = 1.35\%$ through this factor alone.

The whitepaper states its hope in these terms: luminosity uncertainties of about 1% at central rapidity and at scales near the Z mass, and nominal uncertainties that may fall to 0.3–0.5% within a decade (Sect. I C). By Eq. (3.1.8) with ρ near 1, a 1% luminosity needs the gluon to about 0.5% across the range of Figure 2.

Equation (3.1.8) is an estimate. Beyond lowest order $\hat{\sigma}$ is not a delta function, other channels open, and r and ρ are functions of x. The exact calculation is the subject of the next section.

---

4. From error sets to a cross-section

This section gives the rules by which an observable inherits an uncertainty from a PDF set, and the meaning of the correlation ellipses in the whitepaper's Fig. 5.

A set is delivered as a central member with Hessian eigenvector sets, or as an ensemble of Monte Carlo replicas (Sect. V D). Chapter 2.3 derives both. Let σ be any quantity computed from the distributions.

Hessian sets. There are two members for each eigenvector direction k, displaced by one standard deviation each way. With $\sigma_k^{\pm}$ the observable computed from them, linear error propagation gives (derivation D3)

$$\delta\sigma = \frac{1}{2}\sqrt{\sum_k \left(\sigma_k^+ - \sigma_k^-\right)^2}. \tag{3.1.9}$$

Replicas. With $\sigma^{(k)}$ computed from replica k, the prediction and its uncertainty are the mean and standard deviation over the ensemble [2]:

$$\langle\sigma\rangle = \frac{1}{N_{\text{rep}}}\sum_{k=1}^{N_{\text{rep}}}\sigma^{(k)}, \qquad (\delta\sigma)^2 = \frac{1}{N_{\text{rep}}-1}\sum_{k=1}^{N_{\text{rep}}}\left(\sigma^{(k)} - \langle\sigma\rangle\right)^2. \tag{3.1.10}$$

Correlations. Two observables A and B computed from the same set are correlated, with coefficient

$$\rho_{AB} = \frac{1}{4\,\delta A\,\delta B}\sum_k \left(A_k^+ - A_k^-\right)\left(B_k^+ - B_k^-\right) \qquad \text{or} \qquad \rho_{AB} = \frac{\operatorname{cov}(A,B)}{\delta A\,\delta B}. \tag{3.1.11}$$

In the second form $\operatorname{cov}(A,B)$ is the covariance over the replicas, normalised by $N_{\text{rep}} - 1$ as in Eq. (3.1.10). This is the ρ of Eq. (3.1.8), with A and B the gluon at two values of x.

Ellipses. The whitepaper's Fig. 5 shows pairs of cross-sections among $W^{\pm}$, Z, $t\bar{t}$, H and $t\bar{t}H$ at 14 TeV as ellipses. In the Gaussian approximation the joint distribution of (A,B) has contours

$$\begin{pmatrix} a & b \end{pmatrix} \begin{pmatrix} (\delta A)^2 & \rho_{AB}\,\delta A\,\delta B \\ \rho_{AB}\,\delta A\,\delta B & (\delta B)^2 \end{pmatrix}^{-1} \begin{pmatrix} a \\ b \end{pmatrix} = \Delta\chi^2, \qquad P = 1 - e^{-\Delta\chi^2/2}, \tag{3.1.12}$$

with a and b the deviations from the central values and P the probability inside. The ellipse with $\Delta\chi^2 = n^2$ projects onto each axis as $\pm n$ standard deviations but contains only 39% of the probability for $n = 1$ and 86% for $n = 2$. To contain 68.27% or 95.45% it must be drawn at $\Delta\chi^2 = 2.30$ or 6.18. The whitepaper describes its ellipses as rescaled to the 2σ probability; which convention is in use must be checked paper by paper.

The error sets approximate the probability distribution of the fit. They reproduce expectation values and the main correlations and neglect subleading features (Sect. V D).

5. The current sets and how well they agree

This section summarises the whitepaper's comparison of the NNLO sets of 2022 and its assessment of whether their precision can be trusted.

The comparison (Sect. II A) covers NNPDF4.0, CT18, MSHT20, ABMP16 at $\alpha_S(m_Z) = 0.118$, and ATLASpdf21, at a scale of 100 GeV.

DISTRIBUTION	AGREEMENT	THE WHITEPAPER'S EXPLANATION OF THE DIFFERENCES
Gluon	within 1σ for $10^{-4} \lesssim x \lesssim 10^{-1}$	at higher x, different selections and treatments of collider data
Singlet	mostly within 1σ	ABMP16 departs near $x \sim 10^{-3}$; ATLASpdf21 cuts HERA data below $Q^2 = 10\ \mathrm{GeV}^2$
Strangeness	MSHT20 and NNPDF4.0 agree; CT18 lower	the ATLAS 7 TeV W, Z data raise it; CT18 omits them from its default fit because of their high χ^2
Charm	NNPDF4.0 differs from the rest	NNPDF4.0 fits the charm distribution; the others generate it perturbatively
Up valence	agree from $x \sim 0.5$ down to 10^{-2}	below that, weak constraints, with shapes possibly set through the sum rules
Down valence	NNPDF4.0, MSHT20, ABMP16 agree	CT18 and ATLASpdf21 differ, for CT18 partly compensating its lower strangeness

Three causes recur. A single data set, the ATLAS 7 TeV W and Z measurement, separates the groups on strangeness. A choice of method, whether charm is fitted, separates them on charm. And where data are weak, the sum rules of chapter 2.1 carry the choice of functional form into the result.

The uncertainties differ as well. ATLASpdf21 has larger ones because it fits fewer data, ABMP16 may have smaller ones because it defines the 1σ band by $\Delta\chi^2 = 1$, and NNPDF4.0 generally has the smallest luminosity uncertainties.

For luminosities at 14 TeV, all sets agree within uncertainties near $m_X \sim 100$ GeV. Near 1 TeV the gluon–gluon and gluon–quark luminosities of NNPDF4.0 are below those of MSHT20, CT18 and ATLASpdf21 and above ABMP16. NNPDF4.0 includes data sensitive to the large-x gluon that the others do not: dijets at 7 TeV and top-pair distributions from Run II. By section 3, a luminosity at 1 TeV is dominated by x near 0.07 and above.

Precision and accuracy

The whitepaper is direct about the limits of these comparisons (Sect. I C).

- **The data are in tension.** There are disagreements among the most precise experiments, sometimes stronger than random fluctuations would produce. The fit quality per point is elevated in a similar way in CT18, MSHT20 and NNPDF3.1.1.
- **Method matters as much as data.** Differences from fitting technique, treatment of systematics, functional form or definition of the error can be as large as the uncertainty propagated from the data, or larger.
- **Releases move.** The shift from one release of a set to the next can exceed the nominal uncertainty.

The whitepaper asks whether distributions can be trusted at the 1% level and answers with a list of tests. They are closure tests, in which a fit to pseudo-data generated from a known law must reproduce the expected statistics (chapter 2.5); strong goodness-of-fit criteria, which demand internal consistency of the probability distribution in addition to a good global χ^2 [8]; and studies of parametrisation dependence, such as the more than 250 alternative functional forms examined for CT18.

---

6. The data frontier

This section sets out what the LHC data deliver, why more of the same will not be enough, and what the planned facilities are projected to add.

The LHC now

Chapter 1.6 assigns processes to distributions. The whitepaper adds the experimental precision (Sect. III A).

MEASUREMENT	QUOTED PRECISION	MAINLY CONSTRAINS
W and Z/γ^*, differential	about 0.5% at the mass peak	valence quarks, flavour structure of the sea
Inclusive jets and dijets	about 5 to 50%	gluon at large x
W or Z + jets, Z at high p_T	about 15%	gluon and quarks at higher x
Top pairs, total and differential	about 15%	gluon at large x, less strongly than jets
Direct photons, 8 and 13 TeV as ratios	about 5%	gluon at large x, less strongly still

All are compared with NNLO QCD and NLO electroweak predictions.

Why statistics will not be enough

Many of these measurements are already limited by systematic errors, so improvement is not a matter of accumulating data. The reason is one line. Take N points, each with an independent statistical error σ_{stat} and a common systematic shift of standard deviation σ_{sys}. Their average has

$$\text{var} = \frac{\sigma_{\text{stat}}^2}{N} + \sigma_{\text{sys}}^2, \tag{3.1.13}$$

because the common shift does not average. With $\sigma_{\text{stat}} = 5\%$, $N = 50$ and $\sigma_{\text{sys}} = 1\%$ the average is known to 1.22%. A hundred times more data gives 1.00%.

The correlation model

What matters next is how a systematic error is correlated between points (Sect. V A). A source is conventionally taken as fully correlated across all points it affects. Some of the largest are “two-point” systematics, the difference between two Monte Carlo models, and there is little reason for those to be Gaussian or fully correlated.

The cost of a wrong assumption can be computed. Suppose two spectra each prefer their own value $\hat{\lambda}_{1,2}$ of a nuisance parameter λ, with $\chi^2_a(\lambda) = \chi^2_{a,\min} + w_a(\lambda - \hat{\lambda}_a)^2$. Forcing one λ on both and minimising gives

$$\chi^2_{\text{joint}} - \left(\chi^2_{1,\min} + \chi^2_{2,\min}\right) = \frac{w_1 w_2}{w_1 + w_2}\left(\hat{\lambda}_1 - \hat{\lambda}_2\right)^2. \tag{3.1.14}$$

The whitepaper’s example is the ATLAS 8 TeV top-pair data. Separate fits to the transverse-momentum and mass spectra give χ^2 values that add to 11.3. A joint fit with all sources fully correlated gives 45 for 15 points. The parton-shower nuisance parameter is pulled to -1.32 by one spectrum and $+0.39$ by the other, a difference of 1.71 standard deviations. Decorrelating that source alone lowers the joint χ^2 to 14.1, so it accounts for 30.9 units. If the two spectra constrain it equally, Eq. (3.1.14) gives $w = 21$: each spectrum fixes the parameter to ± 0.22 of its assumed width. The two spectra are asked to agree on a modelling difference far more precisely than it was known beforehand.

Two further results from the same section:

- For ATLAS 8 TeV inclusive jets, χ^2 falls from 289 to 248 for 171 points between full correlation and the preferred decorrelation, while the gluon changes little.
- For the four top-pair distributions in the MSHT20 fit, χ^2 per point falls from 6.84 to 1.69 with decorrelation between distributions and to 1.04 with decorrelation within them. Here the change in the gluon is comparable to the difference between an NLO and an NNLO fit.

The effect on the distributions is not large in the best-known region $0.01 < x < 0.1$, but large enough to matter if 1% is the goal.

Projections for the high-luminosity LHC

A study with pseudo-data for the final LHC data set [4] finds that luminosity uncertainties fall by a factor between 2 and 5, depending on the parton combination and the mass (Sect. II B 1). The whitepaper’s headline is a factor of 2 to 3 (Sect. I C, XI), on the condition that the tensions of section 5 are resolved.

Such projections are made by *profiling*: an existing Hessian set is updated with new data without refitting. Let b_k be the displacement along eigenvector k in units of its uncertainty, β_{ik} the change it induces in the prediction T_i for data point i with error σ_i, and T^2 the tolerance, the rise in χ^2 that defines one standard deviation (chapter 2.3). The profiled χ^2 is

$$\chi^2(b) = \sum_i \frac{\left(D_i - T_i - \sum_k \beta_{ik} b_k\right)^2}{\sigma_i^2} + T^2 \sum_k b_k^2. \tag{3.1.15}$$

For one eigenvector and pseudo-data placed at the central prediction, the uncertainty shrinks by (derivation D4)

$$\frac{\delta_{\text{new}}}{\delta_{\text{old}}} = \frac{1}{\sqrt{1 + S^2/T^2}}, \qquad S^2 = \sum_i \frac{\beta_i^2}{\sigma_i^2}. \tag{3.1.16}$$

S measures how many standard deviations of the new data one error set spans. Halving an uncertainty needs $S^2/T^2 = 3$, and a factor of 5 needs $S^2/T^2 = 24$.

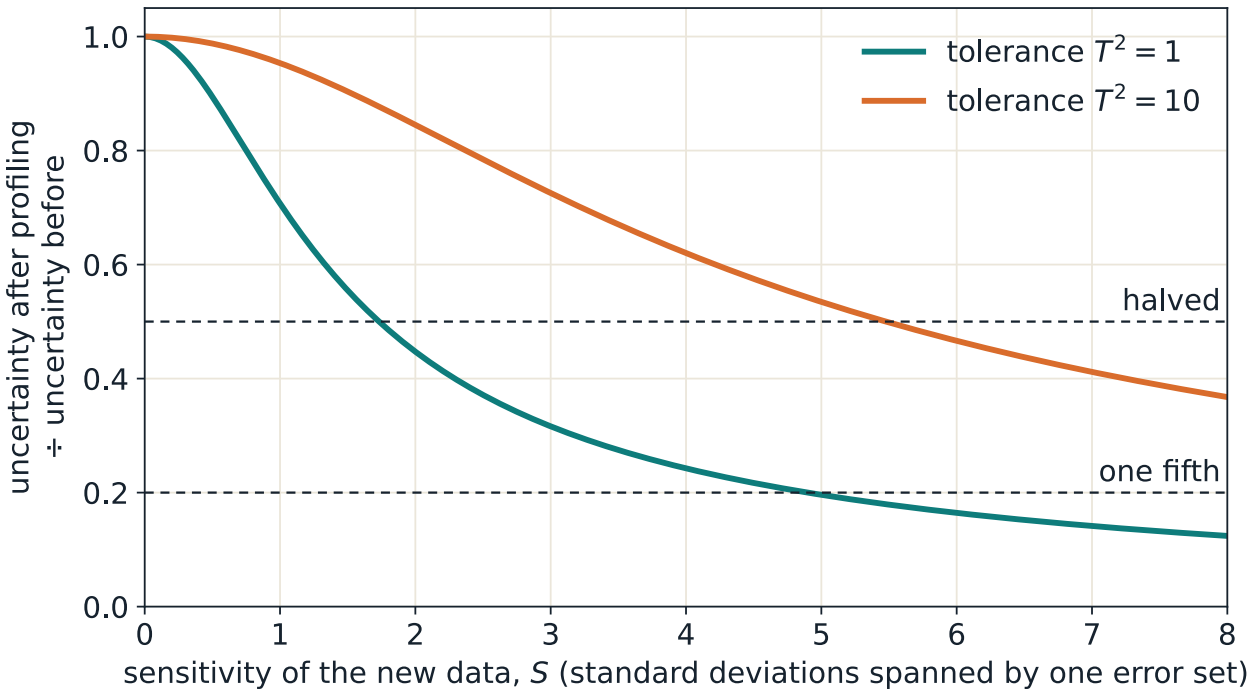

Figure 3. Equation (3.1.16). The same new data shrink an uncertainty much more if the set being profiled defines its errors with $T^2 = 1$ than with a larger tolerance; $T^2 = 10$ is shown as an example. At $S^2 = 3$ the first is halved and the second falls by 12%.

The whitepaper warns that profiling with too aggressive a tolerance overestimates the constraint: the default $\Delta\chi^2 = 1$ of the public profiling code is consistent with the HERAPDF and ATLAS sets and not with CT or MSHT (Sect. IX B 1). A projected reduction factor is meaningful only together with its tolerance.

Lepton–hadron colliders and forward experiments

FACILITY	REACH	PROJECTED EFFECT, AS QUOTED
Electron-Ion Collider (Sect. III B)	$\sqrt{s}$ up to 140 GeV; x up to 0.6–0.7; systematic errors possibly below 1%	valence uncertainties at large x down by up to 80%; sea at small x by up to 50%
LHeC (Sect. III C)	$\sqrt{s} = 1.3$ TeV; $10^{-6} \lesssim x \lesssim 0.9$; total errors 0.8–1.4% in the bulk	luminosity uncertainties down by an order of magnitude; PDF uncertainty on m_W of ± 1.6 MeV against ± 3.7 to ± 5.8 MeV from the high-luminosity LHC
Forward Physics Facility (Sect. III D, III E)	neutrinos from forward hadron production; x down to about 10^{-8}	a region no experiment has probed; charm-tagged neutrino DIS

The Electron-Ion Collider is approved and could record collisions as early as 2030. The LHeC is a proposal. The whitepaper gives two reasons why DIS matters in an era of abundant LHC data (Sect. III C). With one hadron in the initial state, DIS determines each x separately, free of the correlation between small and large x that Eq. (3.1.6) builds into every hadron-collider observable. And it is independent of the systematic limits of the LHC data.

7. The theory frontier

This section states what an N³LO fit requires, what was available in 2022, and what the electroweak sector adds.

Counting orders

The whitepaper (Sect. IV A) expands the splitting functions P and the partonic cross-section c of an observable in $a_s = \alpha_S/4\pi$:

$$\begin{aligned} P &= a_s P^{(0)} + a_s^2 P^{(1)} + a_s^3 P^{(2)} + a_s^4 P^{(3)} + \dots, \\ c &= a_s^n \Big[c^{(0)} + a_s c^{(1)} + a_s^2 c^{(2)} + a_s^3 c^{(3)} + \dots \Big]. \end{aligned} \tag{3.1.17}$$

A fit at N^{k}LO uses the first $k+1$ terms of both. Chapter 1.4 expands in $\alpha_S/2\pi$, so $P^{(k)}$ here is 2^{k+1} times the $P^{(k)}$ of chapter 1.4. NNLO is the present standard.

What N³LO needs

INGREDIENT	STATUS IN THE WHITEPAPER
Four-loop splitting functions $P^{(3)}$	work ongoing
DIS coefficient functions, massless and massive	known
Matching conditions at heavy-quark thresholds	computed almost completely
Evolution codes	several at NNLO, none yet at N³LO
Hadron-collider cross-sections	key inclusive ones, and some differential
Fast interpolation grids at N³LO	to be developed

An approximate N³LO analysis had appeared [9]. The whitepaper notes that more than ten years passed between the NNLO evolution equations and NNLO sets of benchmarked accuracy (Sect. XI).

Where N³LO cross-sections exist, the corrections are a few per cent and the PDF uncertainty is often the largest that remains (Sect. IV B). Until N³LO sets exist, N³LO cross-sections are folded with NNLO sets, at the price of Eq. (3.1.1). Not every process will be known at N³LO, so an N³LO fit will mix orders, and that requires the theory uncertainties of section 8.

The photon and electroweak corrections

At 1% precision, corrections of order α cannot be left out (Sect. IV C). Adding QED splittings to the DGLAP equations makes the photon a parton. The leading term is the abelian copy of gluon radiation from a quark in chapter 1.4, with the colour factor C_F replaced by the squared charge:

$$\begin{aligned} \mu^2 \frac{\partial \gamma(x,\mu^2)}{\partial \mu^2} &= \frac{\alpha}{2\pi} \sum_q e_q^2 \int_x^1 \frac{d\xi}{\xi} P_{\gamma q}\Big(\frac{x}{\xi}\Big) \big[q(\xi,\mu^2) + \bar{q}(\xi,\mu^2) \big] + \dots, \\ P_{\gamma q}(x) &= \frac{1+(1-x)^2}{x}. \end{aligned} \tag{3.1.18}$$

Its size follows from the second moment, $\int_0^1 dx\, x\, P_{\gamma q} = \frac{4}{3}$. Since $e_q^2 \le \frac{4}{9}$ and the quarks carry less than all of the momentum, the photon gains momentum at a rate below $\frac{\alpha}{2\pi}\cdot\frac{4}{3}\cdot\frac{4}{9} = 6.9\times10^{-4}$ per unit of $\ln\mu^2$. From $Q_0 = 1.65$ GeV to M_Z that is

$$\int_0^1 dx\, x\,\gamma(x,M_Z^2)\Big|_{\text{generated above } Q_0} < 6.9\times10^{-4}\times\ln\frac{M_Z^2}{Q_0^2} = 0.55\%. \tag{3.1.19}$$

This is a bound, not an estimate. The photon takes a fraction of a per cent of the momentum from the quarks, which is the size of the precision being sought.

The difficulty was the input at the starting scale. Early sets generated the photon from a model or fitted it to data with little sensitivity, giving uncertainties of 10% or more. The LUXqed approach [7] relates the photon to the measured structure functions of the proton, with elastic and inelastic parts,

$$\gamma(x,Q^2) = \gamma^{\text{el}}(x,Q^2) + \gamma^{\text{inel}}(x,Q^2), \tag{3.1.20}$$

which brings the uncertainty to a few per cent. NNPDF, MMHT, CT and MSHT now provide such a photon, broadly consistent with one another.

Electroweak corrections can be large in the tails. For the total electroweak correction the whitepaper quotes up to about 10% at the highest jet p_T, where the weak part dominates, up to about 20% for the Z at high p_T, and for the precise W and Z data about 0.5% at intermediate and high mass and about 6% in the lowest-mass region. The 0.5% is the size of the experimental error in section 6.

Beyond fixed order

Two further items are listed (Sect. IV D, IV F). Logarithms of $1-x$ and of $1/x$ need resummation at the edges of the kinematic range; a fit with small-x resummation found significant evidence for it in the HERA data (chapter 3.3). And nearly all hadron-collider observables in fits are computed with massless heavy quarks, which will not suffice as precision rises.

---

8. Theory uncertainties in a fit

This section explains how missing higher orders enter a fit, why including them moves central values, and why the result must be used with care in a prediction (Sect. V B).

The theory covariance matrix

Chapter 2.2 showed that a Gaussian theory error, independent of the experimental one, adds to the covariance:

$$\chi^2 = \big(D-T\big)^{\mathsf{T}}\big(C+S\big)^{-1}\big(D-T\big), \qquad S_{ij} = \langle \Delta_i \Delta_j \rangle, \tag{3.1.21}$$

with Δ_i the unknown shift of prediction i. The method was first applied to nuclear corrections, where S is determined empirically, and is part of NNPDF4.0. For missing higher orders S is purely theoretical: the Δ_i are taken from variations of the renormalisation and factorisation scales, in schemes with 5, 7 or 9 points, correlated between processes [6].

The size of such variations is illustrated by the whitepaper's Fig. 22 (Sect. IX C 1). For ATLAS inclusive jets at 8 TeV with $290 < p_T < 318$ GeV, the full range of the cross-section under variation of both scales is about 1.6% at NNLO and 11% at NLO.

Why the central values move

When the method was applied to the NLO NNPDF3.1 fit, the uncertainties grew only a little and the central values moved towards the NNLO result. Both effects appear in the smallest possible example (derivation D5). Two data sets measure one quantity θ with the same experimental error σ. The theory for the second is off by an unknown Δ, to which a variance s^2 is assigned. The best estimate weights the second set by

$$w_2 = \frac{\sigma^2}{2\sigma^2 + s^2} \quad \text{in place of} \quad \frac{1}{2}, \qquad \text{bias} = w_2\,\Delta, \qquad \text{var}(\hat{\theta}) = \frac{\sigma^2(\sigma^2 + s^2)}{2\sigma^2 + s^2}. \tag{3.1.22}$$

With $s = 2\sigma$ the weight falls from $\frac{1}{2}$ to $\frac{1}{6}$. The bias from the faulty theory falls by a factor of 3, while the quoted uncertainty rises only from 0.71σ to 0.91σ. The fit has been rebalanced towards the data whose theory is trusted, which is what brings an NLO fit closer to NNLO.

Correlation with the prediction

The fitted distributions are then used in a prediction with its own missing orders. Take a distribution fixed by one measurement with experimental error σ and theory error s, used to predict an observable whose theory error is also s, with correlation ρ between the two. The error of the prediction is (derivation D5)

$$\text{var} = \sigma^2 + 2s^2\,(1 - \rho). \tag{3.1.23}$$

Treating the two theory errors as independent gives $\sigma^2 + 2s^2$. If they are fully correlated the theory error cancels: a PDF fitted to a process absorbs that process's missing terms and returns them when used to predict the same process. The whitepaper reports that the correlations can often be ignored but in some cases improve both accuracy and precision. Its example is top-pair rapidity distributions predicted from a fit that includes the top-pair total cross-section.

Other routes

- **Sampling the scales.** Each Monte Carlo replica is fitted with its own choice of scales. Using the same scales in the prediction keeps the full correlation.
- **Direct addition.** ATLAS added fully correlated 5-point scale variations of the W and Z predictions to the experimental uncertainty in an NNLO fit.
- **Without scale variation.** Other estimators of missing orders give non-Gaussian priors, which Eq. (3.1.21) cannot hold.

A fit with theory uncertainties can include processes known only at lower order, weighted down by Eq. (3.1.22). The whitepaper expects this to be the route to N³LO fits.

---

9. Benchmarking and the PDF4LHC21 combination

This section describes how the groups compare their fits, how three sets are merged into one, and what the merged uncertainty means (Sect. X, V D).

Benchmarking

The PDF4LHC working group released a combination, PDF4LHC15, in 2015. The 2021 combination [3] uses CT18, MSHT20 and NNPDF3.1.

The central exercise was to fit a reduced, identical data set with theory settings matched as closely as possible. Agreement between the three reduced fits is better than between the global fits, so much of the spread between global fits comes from the choice of data. Differences from the parametrisation, the tolerance and the fitting method remain when the data are the same, in some cases as large as the uncertainties or larger.

The combination

Each set is represented by $N_{\mathrm{rep}} = 300$ replicas, and the 900 are pooled. The mean and variance of the pool are (derivation D6)

$$\bar{f} = \frac{1}{G}\sum_{g=1}^{G} \bar{f}_g, \qquad \delta^2 = \underbrace{\frac{1}{G}\sum_{g=1}^{G} \delta_g^2}_{\text{within}} + \underbrace{\frac{1}{G}\sum_{g=1}^{G} \left(\bar{f}_g - \bar{f}\right)^2}_{\text{between}}, \tag{3.1.24}$$

for G groups with central values $\bar{f}_g$ and uncertainties δ_g. As an example, three determinations of 1.00 ± 0.020, 1.03 ± 0.035 and 0.98 ± 0.015 combine to an uncertainty of 0.032: 0.025 from within and 0.021 from between, in quadrature.

Two properties follow. Disagreement between groups is counted as uncertainty. And the combined uncertainty is never below the root-mean-square of the individual ones, however well the groups agree. That is the behaviour of fully correlated determinations. For two determinations with equal uncertainty δ and correlation ρ, the average has

$$\mathrm{var} = \frac{\delta^2}{2}\,(1 + \rho), \tag{3.1.25}$$

which is δ^2 at $\rho = 1$ and half that for independent results. Pooling assumes $\rho = 1$, the cautious choice for groups that fit largely the same data.

The whitepaper reports a first measurement of ρ. Two methodologies, NNPDF3.1 and NNPDF4.0, fitted to identical data replicas with common theory, give sets that are only partially correlated, because the uncertainty arising from the method is not shared.

Delivery and use

The 900 replicas are reduced to two public sets: 40 Hessian eigenvectors, built by expressing every replica in a common polynomial basis and keeping the principal components, and 100 replicas selected to reproduce the statistics of the pool. PDF4LHC21 agrees with the 2015 combination within uncertainties and is moderately more precise.

The combined sets are intended for searches and simulations. For the most precise measurements, such as electroweak parameters and Higgs physics, comparison with the individual sets remains necessary (Sect. I C).

---

10. Polarised and nuclear distributions

Polarised

In a proton polarised along its motion, a polarised distribution is the difference $\Delta f(x) = f^{\uparrow}(x) - f^{\downarrow}(x)$ between partons with spin along and against the proton's. Its first moment is that parton's contribution to the proton's spin. With $\Delta\Sigma$ the first moment of the quark singlet, ΔG that of the gluon and L the orbital angular momentum,

$$\frac{1}{2} = \frac{1}{2}\,\Delta\Sigma + \Delta G + L. \tag{3.1.26}$$

The Electron-Ion Collider will polarise both beams (Sect. III B 2). The parity-conserving double-spin asymmetry measures the sum of polarised quark and antiquark distributions. The parity-violating one, from the interference of photon exchange with the axial part of Z exchange, measures their difference. The coverage extends to $x \sim 10^{-4}$ and $Q^2 \sim 1000\ \mathrm{GeV}^2$, and the polarised gluon is constrained by scaling violations, as in chapter 1.4.

The projected gains depend on an assumption. With SU(3) flavour symmetry imposed, the uncertainty on ΔG falls by up to 80–90% and that on $\Delta\Sigma$ by about 80%. Without it, the first falls by about 60% and the second shows no clear reduction.

Nuclear

Nuclear PDFs describe a nucleon bound in a nucleus of mass number A, within the same factorisation (Sect. VII). They add A as a variable, with a data set comparable in size to the proton's, so fits assume a smooth dependence on A. That is doubtful for light nuclei such as the deuteron.

Proton and nuclear fits depend on each other. Proton fits use data on nuclear targets, such as neutrino scattering on iron and DIS on the deuteron, to separate flavours, and need nuclear corrections to do so. Nuclear fits take the proton as the boundary at $A = 1$. The whitepaper suggests that simultaneous fits may be needed once the Electron-Ion Collider measures both in one detector, and projects a reduction of nuclear PDF uncertainties by up to a factor of two (Sect. III B 3).

At small x the nuclear structure functions are depleted relative to free nucleons, an effect called shadowing. The whitepaper quotes a 5–15% suppression of neutrino cross-sections at PeV energies (Sect. III D). Saturation effects grow as $A^{1/3}$, and where they set in, collinear factorisation ceases to apply (Sect. VII C).

---

11. Open problems, and where this project sits

The whitepaper's list

Drawn from Table I, Sect. I C, IV F and XI:

1. **Correlated systematic errors.** Better models, built with the experiments (section 6).
2. **Tensions between precise data sets.** Experiments probing the PDFs should reach better agreement among themselves than has been possible so far. The projections assume this.
3. **Theory at N³LO.** Splitting functions, evolution codes, cross-sections and fast interfaces (section 7).
4. **Theory uncertainties** in all fits, with their correlations (section 8).
5. **Large x without nuclear targets** (sections 2 and 10).
6. **New physics in the fitted data** (chapter 3.2).
7. **Small x, heavy quarks and electroweak effects** beyond present approximations (chapter 3.3).
8. **Methodology.** The dependence on functional form, the definition of PDF errors, representative sampling of the space of solutions, and the combination of sets (sections 5 and 9).

This project

This project addresses item 8. The benchmark of section 9 found that method-driven differences survive when the data are common. It could not separate their causes, because the groups differ in parametrisation, tolerance and

fitting technique at once. Our fits remove all but one of these. The data (82 data sets, 4,615 points), the theory (NNPDF theory 40000000: NNLO, $\alpha_S(M_Z) = 0.118$, $Q_0 = 1.65$ GeV) and the parametrisation (the MSHT20 form with 52 parameters) are fixed, in one code, Colibri. Only the method of assigning an uncertainty changes: Hessian, Monte Carlo replicas or Bayesian sampling (chapters 2.3 and 3.4).

Three ledger results connect to this chapter.

- **Fit quality.** On the real global data the fit has $\chi^2/N_{\text{dat}} = 1.2380$ ($\chi^2 = 5713.23$ for 4,615 points, PPDF-35). The ledger records 1.17 for the published NNPDF4.0 fit, a ratio of 1.058. A value above 1 is the pattern the whitepaper reports for the global fits (section 5).
- **Unconstrained directions.** In the closure-test analysis of PPDF-32, 8 of the 52 parameter directions remain unconstrained by the full global data set (chapter 3.4). The whitepaper's concern about representative sampling applies in these directions, and the three methods have room to differ there.
- **Scope.** The loss function of chapter 2.2 contains the experimental covariance and the positivity and integrability penalties. It has no theory covariance, so items 3 and 4 are outside what our uncertainties describe.

Derivations

D1. The cross-section as an integral over the luminosity

Derives (3.1.4), (3.1.5), (3.1.6)

Start from the factorisation theorem of chapter 1.6 for a total cross-section, with a partonic cross-section that depends only on $\hat{s} = x_1 x_2 s$:

$$\sigma = \sum_{i,j} \int_0^1 dx_1 \int_0^1 dx_2 \; f_i(x_1)\, f_j(x_2)\, \hat{\sigma}_{ij}(x_1 x_2 s).$$

The sum is over ordered pairs: i comes from the first proton and j from the second.

Fix the invariant mass. Insert $1 = \int_0^1 d\tau\, \delta(\tau - x_1 x_2)$ and do the x_2 integral, which sets $x_2 = \tau/x_1$ and gives a factor $1/x_1$. The condition $x_2 \le 1$ requires $x_1 \ge \tau$:

$$\sigma = \sum_{i,j} \int_0^1 d\tau\, \hat{\sigma}_{ij}(\tau s) \int_\tau^1 \frac{dx}{x}\, f_i(x)\, f_j\Big(\frac{\tau}{x}\Big).$$

Unordered pairs. For identical hadrons, $\hat{\sigma}_{ij} = \hat{\sigma}_{ji}$. For $i \neq j$ the terms (i,j) and (j,i) together give $\hat{\sigma}_{ij}$ times $\int \frac{dx}{x}\big[f_i(x) f_j(\tau/x) + f_j(x) f_i(\tau/x)\big]$, which is $s\,\mathcal{L}_{ij}$ by Eq. (3.1.2). For $i = j$ there is one term, $\int \frac{dx}{x} f_i(x) f_i(\tau/x)$. The bracket of Eq. (3.1.2) is twice this, and the factor $1/(1+\delta_{ij}) = \frac{1}{2}$ restores it.

Change variable. With $m_X^2 = \tau s$, $s\, d\tau = dm_X^2$, and Eq. (3.1.4) follows. For $\hat{\sigma} = \sigma_0 m_X^2\, \delta(\hat{s} - m_X^2)$ the integral gives Eq. (3.1.5).

Rapidity. With $x = \sqrt{\tau} e^{y}$, $dx/x = dy$ and $\tau/x = \sqrt{\tau} e^{-y}$. The limits $x = \tau$ and $x = 1$ become $y = \mp \ln(1/\sqrt{\tau})$. This is Eq. (3.1.6).

D2. The uncertainty of a luminosity

Derives (3.1.7), (3.1.8)

First order. Replace $g(x)$ by $g(x)\,[1 + r(x)]$ in Eq. (3.1.6):

$$s\,(\mathcal{L}_{gg} + \delta\mathcal{L}_{gg}) = \int dy\; g(x_1)\,g(x_2)\,\big[1 + r(x_1)\big]\big[1 + r(x_2)\big].$$

Dropping the term $r(x_1)r(x_2)$ and dividing by $s\,\mathcal{L}_{gg}$ gives Eq. (3.1.7). The weight $w(y)$ is positive and integrates to 1.

The variance. Treat $r(x)$ as a random function with zero mean. If the weight is concentrated on one pair (x_1, x_2), the variance of Eq. (3.1.7) is that of $r(x_1) + r(x_2)$:

$$\Big(\frac{\delta\mathcal{L}_{gg}}{\mathcal{L}_{gg}}\Big)^2 = \langle r_1^2\rangle + \langle r_2^2\rangle + 2\langle r_1 r_2\rangle = r^2 + r^2 + 2\rho\, r^2,$$

which is Eq. (3.1.8). The general case is an average of such terms over pairs of rapidities.

What was dropped. The second-order term is of relative size r. For $r = 1\%$ and $\rho = -1$ a simulation in the check script gives 1.4×10^{-4} where the first-order formula gives zero.

D3. Error sets, correlations and ellipses

Derives (3.1.9), (3.1.11), (3.1.12)

Hessian sets. Near the best fit, write the parameters as $\theta = \theta_0 + \sum_k z_k\, e_k$, where the e_k are the eigenvectors of the parameter covariance matrix scaled by the square roots of its eigenvalues. The coordinates z_k are uncorrelated with unit variance, and the error sets are the points $z_k = \pm 1$. Linearise an observable: $\sigma(\theta) = \sigma_0 + \sum_k c_k z_k$. Then $\sigma_k^{\pm} = \sigma_0 \pm c_k$, so $c_k = \tfrac{1}{2}(\sigma_k^{+} - \sigma_k^{-})$, and

$$(\delta\sigma)^2 = \Big\langle \Big(\sum_k c_k z_k\Big)^2 \Big\rangle = \sum_k c_k^2 = \frac{1}{4}\sum_k \big(\sigma_k^{+} - \sigma_k^{-}\big)^2.$$

This is Eq. (3.1.9). For two observables the covariance is $\sum_k c_k^A c_k^B$, which gives the first form of Eq. (3.1.11). The formula assumes the observable is linear over the range of the error sets.

The ellipse. Let $v = (a, b)$ be Gaussian with covariance matrix V. The quantity $u = v^{\mathsf T} V^{-1} v$ is the sum of the squares of two independent unit Gaussians. In polar coordinates the probability of $u < \Delta\chi^2$ is

$$P = \int_0^{\sqrt{\Delta\chi^2}} e^{-R^2/2}\, R\, dR = 1 - e^{-\Delta\chi^2/2}.$$

Hence $\Delta\chi^2 = -2\ln(1 - P)$: 2.30 for $P = 0.6827$ and 6.18 for $P = 0.9545$. The extreme value of a on the contour is $\sqrt{\Delta\chi^2\, V_{aa}}$, so the contour $\Delta\chi^2 = n^2$ projects to $\pm n\,\delta A$ and contains $1 - e^{-n^2/2}$: 0.393 for $n = 1$, 0.865 for $n = 2$.

D4. Profiling one eigenvector

Derives (3.1.16)

Take a single eigenvector with coordinate b in Eq. (3.1.15), and pseudo-data generated at the central prediction, $D_i = T_i$:

$$\chi^2(b) = b^2 \sum_i \frac{\beta_i^2}{\sigma_i^2} + T^2 b^2 = \left(S^2 + T^2\right) b^2 .$$

Before. Without the new data, $\chi^2 = T^2 b^2$, and one standard deviation, defined by $\Delta\chi^2 = T^2$, is $b = \pm 1$.

After. The minimum is still at $b = 0$. The criterion $\Delta\chi^2 = T^2$ is now met at $b = \pm T/\sqrt{S^2 + T^2}$. The ratio of the two is Eq. (3.1.16).

Reading it. The prior carries T^2 units of information and the new data add S^2. If the set was built with $T^2 = 10$ but is profiled with $T^2 = 1$, the new data are credited with ten times their weight relative to the old.

D5. A theory covariance in the smallest fit

Derives (3.1.22), (3.1.23)

Two data sets. Let $D_1 = \theta + \varepsilon_1$ and $D_2 = \theta + \Delta + \varepsilon_2$, with $\varepsilon_{1,2}$ independent of variance σ^2. The computed prediction for both is θ. In Eq. (3.1.21), $C = \sigma^2\, 1$ and $S = \mathrm{diag}(0, s^2)$, so

$$\chi^2 = \frac{(D_1 - \theta)^2}{\sigma^2} + \frac{(D_2 - \theta)^2}{\sigma^2 + s^2} .$$

Minimising,

$$\hat\theta = (1 - w_2)\, D_1 + w_2\, D_2, \qquad w_2 = \frac{1/(\sigma^2 + s^2)}{1/\sigma^2 + 1/(\sigma^2 + s^2)} = \frac{\sigma^2}{2\sigma^2 + s^2} .$$

For a given Δ the mean of $\hat\theta$ is $\theta + w_2\Delta$. The variance assigned by the fit is the inverse of the sum of the weights, $\left[1/\sigma^2 + 1/(\sigma^2 + s^2)\right]^{-1}$. These are the statements of Eq. (3.1.22).

One data set and a prediction. Let the distribution be fixed by $D = \theta + \Delta + \varepsilon$, so $\hat\theta = D$. The observable to be predicted is truly $\theta + \Delta'$ and is computed as $\hat\theta$. The error of the prediction is

$$\hat\theta - (\theta + \Delta') = \varepsilon + \Delta - \Delta' .$$

With $\langle \Delta^2 \rangle = \langle \Delta'^2 \rangle = s^2$ and $\langle \Delta\Delta' \rangle = \rho s^2$, its variance is $\sigma^2 + 2s^2 - 2\rho s^2$, which is Eq. (3.1.23).

D6. The mean and variance of a pooled ensemble

Derives (3.1.24), (3.1.25)

Pool equal numbers of replicas from G sets. A replica drawn at random from the pool belongs to set g with probability $1/G$. Within set g the quantity of interest has mean $\bar f_g$ and variance δ_g^2, so $\langle f^2 \rangle_g = \delta_g^2 + \bar f_g^2$.

Mean. $\langle f \rangle = \frac{1}{G}\sum_g \bar f_g = \bar f$.

Variance.

$$\langle f^2\rangle - \bar{f}^2 = \frac{1}{G}\sum_g \left(\delta_g^2 + \bar{f}_g^2\right) - \bar{f}^2 = \frac{1}{G}\sum_g \delta_g^2 + \frac{1}{G}\sum_g \left(\bar{f}_g - \bar{f}\right)^2,$$

using $\frac{1}{G}\sum_g \bar{f}_g^2 - \bar{f}^2 = \frac{1}{G}\sum_g (\bar{f}_g - \bar{f})^2$. This is Eq. (3.1.24).

Two correlated determinations. For f_1, f_2 with variance δ^2 each and covariance $\rho\,\delta^2$, the average $\frac{1}{2}(f_1 + f_2)$ has variance $\frac{1}{4}(\delta^2 + \delta^2 + 2\rho\,\delta^2)$, which is Eq. (3.1.25).

References

1. S. Amoroso et al., “Snowmass 2021 whitepaper: Proton structure at the precision frontier”, arXiv:2203.13923. The primary source; its section numbers are cited in the text. References [3] to [9] are quoted from its bibliography.
2. J. M. Moore, *Parton Distributions in Beyond the Standard Model Theories*, PhD thesis, University of Cambridge, sections 1.3.4 and 1.4.
3. R. D. Ball et al. (PDF4LHC Working Group), J. Phys. G 49 (2022) 080501, arXiv:2203.05506. The PDF4LHC21 combination.
4. R. Abdul Khalek, S. Bailey, J. Gao, L. Harland-Lang and J. Rojo, Eur. Phys. J. C 78 (2018) 962, arXiv:1810.03639. Projections for the high-luminosity LHC.
5. C. Anastasiou et al., JHEP 05 (2016) 058, arXiv:1602.00695. The origin of Eq. (3.1.1).
6. R. Abdul Khalek et al. (NNPDF), Eur. Phys. J. C 79 (2019) 838, arXiv:1905.04311, and Eur. Phys. J. C 79 (2019) 931, arXiv:1906.10698; R. D. Ball and R. L. Pearson, Eur. Phys. J. C 81 (2021) 830, arXiv:2105.05114. Missing higher orders as a theory covariance matrix.
7. A. V. Manohar, P. Nason, G. P. Salam and G. Zanderighi, Phys. Rev. Lett. 117 (2016) 242002, arXiv:1607.04266, and JHEP 12 (2017) 046, arXiv:1708.01256. The LUXqed photon.
8. K. Kovařík, P. M. Nadolsky and D. E. Soper, Rev. Mod. Phys. 92 (2020) 045003, arXiv:1905.06957. Strong goodness-of-fit tests.
9. J. McGowan, T. Cridge, L. A. Harland-Lang and R. S. Thorne, Eur. Phys. J. C 83 (2023) 185, arXiv:2207.04739. The approximate N³LO analysis cited by the whitepaper.