

Reconstruction of Monte Carlo replicas from Hessian parton distributions

Tie-Jiun Hou,^a Jun Gao,^{b,c} Joey Huston,^d Pavel Nadolsky,^a Carl Schmidt,^d
Daniel Stump,^d Bo-Ting Wang,^a Ke Ping Xie,^a Sayipjamal Dulat,^{d,e,f} Jon Pumplin^d
and C.P. Yuan^d

^aDepartment of Physics, Southern Methodist University,
Dallas, TX 75275-0181 U.S.A.

^bINPAC, Shanghai Key Laboratory for Particle Physics and Cosmology,
Department of Physics and Astronomy, Shanghai Jiao-Tong University, Shanghai 200240, China

^cHigh Energy Physics Division, Argonne National Laboratory,
Argonne, Illinois, 60439 U.S.A.

^dDepartment of Physics and Astronomy, Michigan State University,
East Lansing, MI 48824 U.S.A.

^eSchool of Physics Science and Technology, Xinjiang University,
Urumqi, Xinjiang 830046 China

^fCenter for Theoretical Physics, Xinjiang University,
Urumqi, Xinjiang 830046 China

E-mail: tiejiunh@mail.smu.edu, jgao@anl.gov, huston@pa.msu.edu,
nadolsky@smu.edu, schmidt@pa.msu.edu, stump@pa.msu.edu,
botingw@mail.smu.edu, kepingx@mail.smu.edu, sdulat@hotmail.com,
pumplin@pa.msu.edu, yuan@pa.msu.edu

ABSTRACT: We explore connections between two common methods for quantifying the uncertainty in parton distribution functions (PDFs), based on the Hessian error matrix and Monte-Carlo sampling. CT14 parton distributions in the Hessian representation are converted into Monte-Carlo replicas by a numerical method that reproduces important properties of CT14 Hessian PDFs: the asymmetry of CT14 uncertainties and positivity of individual parton distributions. The ensembles of CT14 Monte-Carlo replicas constructed this way at NNLO and NLO are suitable for various collider applications, such as cross section reweighting. Master formulas for computation of asymmetric standard deviations in the Monte-Carlo representation are derived. A correction is proposed to address a bias in asymmetric uncertainties introduced by the Taylor series approximation. A numerical program is made available for conversion of Hessian PDFs into Monte-Carlo replicas according to normal, log-normal, and Watt-Thorne sampling procedures.

KEYWORDS: QCD Phenomenology, Deep Inelastic Scattering (Phenomenology)

ARXIV EPRINT: [1607.06066](https://arxiv.org/abs/1607.06066)

Contents

1	Introduction	1
2	Generation of Monte-Carlo replicas from Hessian error PDFs	3
2.1	Master formulas for Hessian PDFs	3
2.2	Probability density distribution for X	7
2.3	Generation of Monte-Carlo replicas	9
2.4	Monte-Carlo replicas for CT14 parton distributions	11
3	Comparisons of replica sets and Hessian error PDFs	12
3.1	Estimation of PDF uncertainties	12
3.2	Only an ensemble of MC replicas is meaningful	17
4	Conclusions	20
A	Definition of χ^2 in CT fits	22

1 Introduction

Modern parton distribution functions (PDFs) [1–6] are provided with estimates of uncertainties from multiple origins. These estimates are essential for understanding of the accuracy of collider predictions, both for precision measurements and for new physics searches. PDFs are determined from a global statistical analysis of diverse experimental measurements, in deep-inelastic scattering, in production of vector bosons and jets, and in other hard-scattering processes. An optimal parametrization of the PDFs is obtained by a minimization of the figure-of-merit function (χ^2) that quantifies the level of agreement between experimental data and theoretical predictions. The minimization is performed with respect to the PDF parameters of interest, and with respect to nuisance parameters associated with theory, experiment and the data analysis procedure. Once the best-fit PDF parametrization is found, additional parametrizations are constructed for estimating the total PDF uncertainty. In this paper, we explore the connection between two methods for quantifying PDF uncertainties, one based on the diagonalization of the Hessian error matrix [7], and one on the stochastic (Monte-Carlo) sampling of parton distributions [8, 9]. PDF uncertainties can also be determined by the Lagrange multiplier [10] and offset [11] methods, but the PDFs obtained with the Hessian and Monte-Carlo techniques are the most commonly used.

In any method of PDF error analysis, the ultimate goal is to provide information about the probability distribution in the space of PDF parameter values. This information can be presented in several forms. Much of PDF research [1, 2, 4–6] relies on an analytic

χ^2 minimization on a class of PDF parametrization forms, in the same manner as in the CT global analyses. Since $\chi^2(\vec{a})$ can be approximated by a quadratic function of PDF parameters \vec{a} in the neighborhood of the global minimum, a boundary of the hyperellipsoid containing a specified cumulative probability p can be delineated with a relatively small number of error PDFs, corresponding to the eigenvectors of the Hessian matrix. The PDF uncertainty on a QCD observable X , at a confidence level p , is derived via algebraic “master formulas” [7, 12] from values $X_{\pm i}$ for X calculated for each eigenvector set i . In the Gaussian approximation, the PDF errors on X in the positive and negative directions are equal. If deviations from the Gaussian behavior are mild, then the asymmetry between the positive and negative PDF errors can be estimated in the Hessian approach, too, using master formulas requiring two error sets per each eigenvector direction [13].

An alternative NNPDF Monte-Carlo technique [3, 14–17] provides an ensemble of error PDF sets, or “replicas”, which sample the functional forms of parton distributions parametrized by neural networks, with little bias due to the choice of the parametrization form. The probability distribution in the hypervolume can in principle be fully reconstructed, given a sufficiently large ensemble of PDF replicas. The prediction for a QCD observable X is then obtained from the mean and standard deviation of the values of X for each member of the (large) Monte-Carlo replica ensemble.

Hessian eigenvector sets can be converted into Monte Carlo replicas [18], and vice versa [19–21]. Small Hessian eigenvector ensembles are sufficient, and desirable, for many applications. Monte Carlo calculations based on a large ensemble of replica PDFs allow the implementation of new experimental constraints on the PDFs using replica re-weighting [8, 22, 23]. The full ensemble of Monte Carlo replicas can be recast into an ensemble with fewer replicas by unweighting [23], by compression [24], or by conversion into a Hessian ensemble [19–21], while retaining the core statistical information. PDF ensembles from several groups can be combined by converting the Hessian ensembles into a Monte-Carlo representation, and then reducing/compressing the combined MC ensemble into a smaller Hessian or MC ensemble. These techniques are employed to combine the error PDFs from CT14, MMHT’14, and NNPDF3.0 into smaller PDF4LHC15 ensembles according to the recommendation of the PDF4LHC Working Group [25].

Conversion of the Hessian eigenvector sets into Monte-Carlo replica sets is employed in several such applications and influences the outcomes for the uncertainties and even central predictions. The PDF conversion must be understood in order to trust this method. The Watt-Thorne prescription [18] provides the simplest realization of the conversion method, under the approximations of the exactly Gaussian distribution of the PDF parameters and linear dependence of error PDFs on small variations of the PDF parameters. It is desirable, however, to develop a formalism that can go beyond the linear approximation and, given enough Hessian sets, reproduce non-linear features of the probability distribution such as asymmetry.

This paper lays out such general formalism. It allows one to systematically include nonlinear PDF dependence and may produce different outcomes in some LHC predictions compared to the Watt-Thorne method, as we will demonstrate. To apply the general conversion method, we first construct Monte-Carlo replica ensembles that reproduce both

the *symmetric* and *asymmetric* uncertainties of the CT14 Hessian sets, and also generate positive-definite replica parametrizations if desired. The resulting PDF uncertainties with MC replicas are closer to CT14 asymmetric uncertainties, compared to the Watt-Thorne method. Reconstruction of the full probability distribution from the Hessian sets reveals several subtleties, see section 2B. In particular, naive Monte-Carlo sampling of the PDFs using the Taylor series results in biased asymmetric uncertainty bands that can be corrected.

It is also instructive to compare statistical properties of the CT14 MC replica ensemble (obtained by conversion) and NNPDF MC replica ensemble (obtained using a genetic algorithm). Despite drastic differences in the two replica generation methods, we demonstrate in section 3B that the statistical properties of two MC replica ensembles have deep similarities, and explain why.

In section 2 we compare formulas for the estimation of Hessian and MC uncertainties and derive prescriptions for the generation of MC replicas that generalize the Watt-Thorne prescription. These prescriptions account for asymmetry of PDF errors in the same manner as in the CT14 Hessian set. At the end of the section, we show how to construct individual replica sets that reproduce the positivity requirement imposed on the CT14 Hessian sets. Section 3 compares PDFs, parton luminosities, predictions for LHC cross sections, and their uncertainties obtained with the CT14 Hessian and MC PDFs. In section 3.2 we examine statistical properties of the replica ensemble, and consider the agreement of individual replicas with the experiments in the global fit. Section 4 contains concluding remarks and information on the availability of the CT14 replica PDFs. A computer program is presented for generation of MC replicas according to the new prescriptions that optimally reproduce the Hessian uncertainties and positivity.

2 Generation of Monte-Carlo replicas from Hessian error PDFs

2.1 Master formulas for Hessian PDFs

The CT14 parton distribution functions (PDFs) were generated in a global analysis of QCD, by fitting theoretical calculations, at LO, NLO and NNLO, to experimental data for a wide variety of high- Q^2 processes [1]. Once the central (most probable) combination of PDF parameters is found by minimization of the log-likelihood function χ^2 , special “error PDF sets” are constructed to characterize and propagate PDF uncertainties. The error PDFs do not specify all features of the probability distribution in the fit, only a part needed for estimating uncertainties in typical applications. In the Hessian method used by CTEQ, the error PDFs, called “eigenvector sets”, keep information about the first and second moments of the probability distribution, sufficient for estimation of central values, standard deviations, and PDF-driven correlations. A PDF ensemble based on MC replicas, when obtained directly from the fit, can in principle reproduce the primordial probability distribution with better accuracy; in practice, its size is commonly limited to no more than a thousand replicas, also enough for reproducing up to the second moments.

Therefore, the error sets primarily quantify the lowest two moments of the primordial probability; and a Hessian eigenvector set is often the only published output from the PDF




Figure 1. Contours of constant χ^2 under the scaling transformation of PDF parameters.

fit. For the applications utilizing MC replicas, especially when large numbers of replicas are to be generated on the fly, one may wish to convert the Hessian PDF ensemble into an MC replica one, according to the procedure that will be laid out. [A computer code implementing this procedure is available for downloading.] The converted MC replicas must ideally preserve the statistical features of the Hessian ensemble. They cannot provide more information about the primordial higher moments than what is available in the Hessian ensemble.

Let us focus for a moment on the construction and properties of the CT Hessian sets. The PDFs $f_n(x, Q)$ are parametrized at $Q = Q_0 = 1.3$ GeV by a functional form $F_n(x, \vec{a})$ which depends on a set of D adjustable parameters, $\vec{a} = \{a_1, a_2, a_3, \dots, a_D\}$. A test statistic $\chi^2(\vec{a})$ is defined as in [26, 27] to measure the difference between theory and data. Its explicit definition is reproduced in the appendix. The minimum of $\chi^2(\vec{a})$ defines the central fit, and the variation of χ^2 in a neighborhood of the minimum provides the uncertainty in the fit. Thus, the approximation in the vicinity of the minimum,

$$\chi^2(\vec{a}) \approx \chi^2(\vec{0}) + \frac{1}{2} \sum_{i,j=1}^D \left(\frac{\partial^2 \chi^2}{\partial a_i \partial a_j} \right)_{\vec{a}=\vec{0}} a_i a_j = \chi_0^2 + \sum_{i,j=1}^D H_{ij} a_i a_j, \quad (2.1)$$

has two parts. First, the central combination of the parameters, which corresponds, without loss of generality, to all zero a_i values ($\vec{a} = \{0, 0, \dots, 0\} \equiv \vec{0}$), gives the “best” fit between theory and data, for which χ^2 is minimized: $\chi^2(\vec{0}) = \min \chi^2 \equiv \chi_0^2$. Second, the approximately quadratic behavior of χ^2 very close to the minimum allows one to define independent directions in \vec{a} space, corresponding to eigenvectors \hat{z}_i of the Hessian matrix H_{ij} . The eigenvectors are used to determine the uncertainty on PDFs f and on any quantities $X(f)$ that depend on them, by computing variations of χ^2 along each independent direction [7].

In this scheme, the independent directions \hat{z}_i are determined by relying on the quadratic approximation for χ^2 ; at the same time, extreme displacements $\bar{z}_{\pm i}$ along each eigenvector direction are found from variations of the true (not perfectly quadratic) χ^2 . This is illustrated in the left inset of figure 1, showing contours of constant values of the exact and approximate χ^2 for a certain pair of parameters, $\{z_i, z_j\}$. The minimum of the exact χ^2 is reached at the axes origin; the solid and dashed contours are drawn for constant increases

of $\Delta\chi_i^2$ with respect to the minima of the exact χ^2 and its quadratic approximation. The \hat{z} directions are along the axes of the approximate χ^2 ellipsoid. The extreme displacements $\bar{z}_{\pm i}$ are found by demanding the increase T^2 in the true, not approximate, χ^2 . They correspond to the (filled) red circles at the intersections of the z axes and the irregular contours of the true χ^2 in the left-hand side of figure 1. Note that both the functional form of the true χ^2 and the value of T^2 depend on the context of the analysis. In CT fits [26, 27], the true χ^2 that determines the uncertainties consists of two tiers of conditions, quantifying the global agreement with the experiments, and deviations from individual experiments, cf. the appendix. The parameter T is set to 10 at 90% confidence level (c.l.). The complex nature of true χ^2 implies that, most generally, $\bar{z}_{+i} \neq \bar{z}_{-i}$; still, its overall features are captured well by the approximate quadratic χ^2 .

This introduces one source of asymmetry in PDF errors, due to the cubic and higher sign-odd powers of z in the full $\chi^2(z)$. The impact of this source is reduced by rescaling $z_i \rightarrow R_i(z_i)$ of the parameters to satisfy $R_i(\vec{0}) = 0$, $R_i(\bar{z}_{\pm i}) \equiv R_{\pm i} = \pm 1$. Keep in mind that disagreements of the primordial z -dependent probability with the Gaussian behavior are weak to start with, and that the Hessian sets specify the probability (i.e., χ^2) only at three points per an eigenvector direction. *Aside from the above conditions at the three points, we will not need to know the rescaling function $z_i(R_i)$ in the rest of the discussion.*¹ The probability distribution in terms of \vec{R} is even closer to the Gaussian one:

$$\mathcal{P}(\vec{R}) \approx \frac{1}{(2\pi)^{D/2}} \exp\left(-\frac{1}{2} \sum_{i=1}^D R_i^2\right). \quad (2.2)$$

The transformation is illustrated by the right-hand side of figure 1. This approximation is known from experience to hold well in CT fits within the 90% c.l. region around the best fit, even though large $|R_i|$ (greater than 1.5) may be excluded by tier-2 penalties on χ^2 that enforce agreement with the individual fitted experiments [26, 27].

Next, we consider a function $X(\vec{R})$ of the parameters R_i , such as the PDF or even the QCD cross section. This function may have asymmetric uncertainties, regardless of symmetry of $\mathcal{P}(\vec{R})$. Expanding $X(\vec{R})$ near the global χ^2 minimum in the Taylor series up to the second derivatives, we have

$$X(\vec{R}) = X_0 + \sum_{i=1}^D \frac{\partial X}{\partial R_i} R_i + \frac{1}{2} \sum_{i,j=1}^D \frac{\partial^2 X}{\partial R_i \partial R_j} R_i R_j + \dots \quad (2.3)$$

The derivatives can be estimated by finite differences,

$$\frac{\partial X}{\partial R_i} \approx \frac{X_{+i} - X_{-i}}{2}, \quad (2.4)$$

$$\frac{\partial^2 X}{\partial R_i^2} \approx X_{+i} + X_{-i} - 2X_0, \quad (2.5)$$

$$\frac{\partial^2 X}{\partial R_i \partial R_j} \approx \frac{X_{+i,+j} + X_{-i,-j} - X_{+i,-j} - X_{-i,+j}}{2}. \quad (2.6)$$

¹An example of the rescaling function for one eigenvector direction is a steplike one, $z_i(R_i) = |R_i/2| [\bar{z}_{+i} + \bar{z}_{-i} + (\bar{z}_{+i} - \bar{z}_{-i}) \tanh(R_i/\epsilon)]$. The parameter ϵ controls the smoothness of the step transition, it is at least as small as $(\bar{z}_{+i} + \bar{z}_{-i})/(\bar{z}_{+i} - \bar{z}_{-i})$. Far enough from $z_i = R_i = 0$, the step function reduces to $z_i = |R_i| \bar{z}_{+i}$ and $z_i = |R_i| \bar{z}_{-i}$, for positive and negative z_i , respectively.

In these equations, the X values are taken at the global minimum,

$$X_0 \equiv X(\vec{0}); \tag{2.7}$$

at the red circles on the axes in figure 1,

$$X_{\pm i} \equiv X(0, 0, \dots, R_i = \pm 1, \dots, 0); \tag{2.8}$$

and at the blue squares on the diagonals,

$$X_{\pm i, \pm j} \equiv X\left(0, \dots, R_i = \pm \frac{1}{\sqrt{2}}, \dots, R_j = \pm \frac{1}{\sqrt{2}}, \dots, 0\right). \tag{2.9}$$

We see that the $2D$ Hessian eigenvector sets corresponding to $R_i = \pm 1$ are sufficient for computing the first and diagonal second derivatives, $\partial X / \partial R_i$ and $\partial^2 X / \partial R_i^2$. The mixed second derivatives would require $2D(D - 1)$ additional error sets for $\{R_i, R_j\} = \{\pm 2^{-1/2}, \pm 2^{-1/2}\}$. Those off-diagonal error sets are not provided in the conventional PDF analyses.

The Taylor expansion of $X(\vec{R})$ can be used to construct various master formulas for the estimation of Hessian PDF uncertainties. Keeping only linear terms on the right-hand side of eq. (2.3), one obtains

$$X(\vec{R}) = X_0 + \sum_{i=1}^D \frac{X_{+i} - X_{-i}}{2} R_i. \tag{2.10}$$

If $R_i = \pm 1$ are reached at the boundary of the 68% c.l. region (and the Hessian eigenvector sets are determined according to this convention), the maximal variation of X within the unit hypersphere of \vec{R} ,

$$\delta_{68}^H X = |\nabla X| = \frac{1}{2} \sqrt{\sum_{i=1}^D [X_{+i} - X_{-i}]^2}, \tag{2.11}$$

yields the Hessian symmetric PDF uncertainty at 68% c.l., denoted by $\delta_{68}^H X$ [7]. However, the published CT14 parametrizations correspond to the 90% c.l. Using the same master formula, we find the 90% c.l. uncertainty from the CT14 sets, satisfying $\delta_{90}^H X \approx 1.65 \delta_{68}^H X$.

If the *diagonal* second-order derivatives are also included, eqs. (2.10) and (2.11) are modified to

$$X(\vec{R}) = X_0 + \sum_{i=1}^D \left(\frac{X_{+i} - X_{-i}}{2} R_i + \frac{X_{+i} + X_{-i} - 2X_0}{2} R_i^2 \right). \tag{2.12}$$

Also, asymmetric uncertainties on X can be estimated by [13]

$$\delta_{68}^{H>} X = \sqrt{\sum_{i=1}^D [\max(X_{+i} - X_0, X_{-i} - X_0, 0)]^2}, \tag{2.13}$$

$$\delta_{68}^{H<} X = -\sqrt{\sum_{i=1}^D [\max(X_0 - X_{+i}, X_0 - X_{-i}, 0)]^2}. \tag{2.14}$$

The absolute values of the asymmetric uncertainties are generally not the same for the positive ($>$) and negative ($<$) displacements [13]. For example, $X_{+i} - X_0$ and $X_{-i} - X_0$ must have opposite signs if X varies linearly with R_i , but they can have the same sign if the nonlinearity of X is large and hence result in a zero contribution to either $\delta_{68}^{H>}X$ or $\delta_{68}^{H<}X$. In the CT convention, low-number eigenvectors, corresponding to the best-determined directions, tend to be more symmetric (with opposite-sign variations), while high-number eigenvectors can be strongly asymmetric.

2.2 Probability density distribution for X

We will now formulate a set of requirements for Monte-Carlo replica generation that will be applied in the next subsection to devise our replica generation method. A reader interested in the specific formulas for replica generation may skip this subsection and proceed to the next.

The Monte-Carlo replicas quantify the full probability density distribution $\mathcal{P}(X)$ for a QCD observable X , going beyond the intervals of fixed probability available with the Hessian PDFs. A priori, $\mathcal{P}(X(\vec{R})) \neq \mathcal{P}(\vec{R})$. The probability $\mathcal{P}(X(\vec{R}))$ might be extracted *directly* from the fit by using a Lagrange multiplier scan or stochastic sampling of the *exact* log-likelihood χ^2 . But, more often than not, the outcomes of the fit are initially stored in the form of the Hessian PDFs $f_a(x, Q; \vec{R})$. The complete functional forms for the PDFs may be unavailable or difficult to sample. Then, one relies on constructing the Monte-Carlo replicas from the Taylor series for the Hessian PDF values provided in $\{x, Q\}$ space, and without invoking the exact parametrization form.

In addition to the central prediction X_0 , the Hessian PDFs specify two values of X per each eigenvector direction, that is, $X_{\pm i}$. This information is generally insufficient for reconstructing the true $\mathcal{P}(X)$ by Monte-Carlo sampling. Even the confidence of probability intervals for $\mathcal{P}(X)$ may remain undetermined!

Recall that, by their construction and using the parameter distribution $\mathcal{P}(\vec{R})$, the 68% c.l. Hessian values $\{X_{-i}, X_0, X_{+i}\}$ correspond to the cumulative probabilities of about 16%, 50%, and 84%. When one naively samples $X(\vec{R})$ assuming a multi-Gaussian distribution for \vec{R} , one is not guaranteed to obtain the same cumulative probabilities for $\{X_{-i}, X_0, X_{+i}\}$ from the distribution $\mathcal{P}(X)$ of Monte-Carlo replicas, because of *probability folding*. When folding occurs, the 68% probability interval of the Monte-Carlo replica set distributed according to $\mathcal{P}(X)$ deviates from the 68% c.l. Hessian uncertainty interval found using $\mathcal{P}(\vec{R})$. In the Gaussian case, the central combination $\vec{R} = \vec{0}$ of the PDF parameters, associated with the central CT14 PDF, is located at the mode, median, and mean of $\mathcal{P}(\vec{R})$, as those coincide. On the other hand, the central CT14 prediction for X does not automatically correspond to either the mode, median, or mean of $\mathcal{P}(X)$, because those can differ.

Indeed, when $X(\vec{R})$ is not monotonic, several combinations $(\vec{R}(X))_k$ may produce the same value of X . Replicas of such X obey a folded probability $\mathcal{P}(X)$, meaning that the cumulative probability for getting $X(\vec{R})$ is not equal to that for the corresponding \vec{R} . In the one-dimensional case, a simple example is $X = |R|$ with a normally distributed R on $[-\infty, +\infty]$. Then X obeys a half-normal distribution on $X \in [0, +\infty]$;

$X_0 = 0$ corresponds to the cumulative probabilities of 50% and 0% according to $\mathcal{P}(R)$ and $\mathcal{P}(X(R))$, respectively.

If $X(R)$ depends on one eigenvector parameter R and is continuous and monotonic in K intervals $[r_0, r_1], \dots, [r_{k-1}, r_k]$, then $X(R)$ is invertible in each interval: $R = R_k(X)$ for $r_{k-1} \leq R \leq r_k$. We indicate that x lies within $[a, b]$ using a notation $\Theta(x \in [a, b])$, with $\Theta(L)$ equal to 1 or 0 if the statement L is true or false, respectively.

The function $X(R)$ is generally not normally distributed, even when R is. That is, while

$$\mathcal{P}(R) = \frac{1}{(2\pi)^{1/2}} e^{-R^2/2}, \tag{2.15}$$

the probability density distribution $\mathcal{P}(X)$ does not need to be Gaussian. It satisfies

$$\mathcal{P}(X) = \sum_{k=1}^K \Theta(R \in [r_{k-1}, r_k]) \Theta(X \in [X(r_{k-1}), X(r_k)]) \mathcal{P}(R_k(X)) \frac{dR_k(X)}{dX}, \tag{2.16}$$

as follows from the normalization conditions for the probability densities,

$$\int \mathcal{P}(R) dR = \int \mathcal{P}(X) dX = 1. \tag{2.17}$$

Going back to the truncated Taylor series with a non-negligible $\partial^2 X / \partial R^2$,

$$X(R) = X_0 + \frac{\partial X}{\partial R} R + \frac{1}{2} \frac{\partial^2 X}{\partial R^2} R^2 + \dots, \tag{2.18}$$

where the derivatives are estimated by the finite differences from X_0 and X_{\pm} , we could solve for $R(X)$ and find up to *two* solutions for R per each X value. [$\mathcal{P}(X)$ is then given by eq. (2.16) with $K = 2$.] Such truncated Taylor approximation is clearly not monotonic. We expect it to deviate from the true $X(R)$ when the displacements R are large, and to produce a folded probability distribution $\mathcal{P}(X)$ with a displaced median value and confidence intervals, as compared to the true $\mathcal{P}(X)$. We may even end up with a problematic arrangement when the central CT14 Hessian prediction is outside of the nominal 68% probability interval of the replicas. These observations continue to hold when \vec{R} has D components.

On the other hand, in a variety of important applications, our group observed that the Lagrange multiplier and Hessian methods render similar probability intervals. Most recently, the two methods were shown to predict compatible 68% and 90% confidence intervals for the common LHC observables, such as jet, Higgs, and $t\bar{t}$ production cross sections [1, 28–31]. This indicates that the impact of the second derivatives is mild in constrained regions, credible estimates for $\mathcal{P}(X)$ can be found via the sampling of the Hessian PDFs after correcting for the Taylor series' imperfections. At intermediate x , where the nonlinearities in the PDF uncertainties are weak, the Taylor expansion performs well. The discrepancies are more pronounced at small and large Bjorken x with fewer experimental constraints. Nonlinearities and folding effects are small for well-constrained eigenvector parameters R_i and can be substantial for poorly constrained ones.

With these observations in mind, we design a Monte-Carlo sampling procedure for the PDFs using the Taylor series in the case of mild nonlinearities. First, rather than always sampling the PDFs $f_a(x, Q; \vec{R})$, we may sample a function X of $f_a(x, Q; \vec{R})$ that reduces the impact of $\partial^2 X / \partial R_i^2$. For example, at moderate x , one could just sample the PDFs directly, i.e., set $X = f_a(x, Q; \vec{R})$ with normally distributed \vec{R} . At large x , where $f_a(x, Q) \sim (1-x)^R$, Gaussian sampling of $X \equiv \ln(f_a(x, Q)) \sim R \cdot \text{const}$ may be more desirable for suppressing the nonlinearities; we will show that the log-normal sampling leads to more physical behavior. We explore this possibility when generating positive-definite PDF replicas in section 2.4.

Second, we need to assign cumulative probabilities to the input X values. Within the accuracy of the Hessian approximation, we can make use of the knowledge that the Hessian central value and Hessian uncertainties are close to the median and 68% central probability intervals for the PDFs, respectively. This “obvious” conclusion follows from the Lagrange Multiplier studies. It can be violated by naive sampling, hence we enforce it as a separate requirement. Furthermore, for mildly asymmetric populations, the median, with the cumulative probability of 0.5, tends to lie close to the mean given by the standard formula [32, 33]. In the various sampling scenarios that we had tried, we confirmed that the median of a PDF replica distribution differed little from the mean, except for in the extreme regions. It suffices to assume that the central Hessian PDF coincides with the *mean* of the PDF replica ensemble, easily computable from the replicas’ values and close to the median value. So when the replica ensemble is first generated, the median/mean at each x and Q may disagree from the central CT14 PDF as a consequence of the probability folding caused by the Taylor expansion, as well as because of residual fluctuations. We correct for this mismatch by a uniform shift in all PDF replica values, as explicated in the next section.

In section 3, we show that this sampling procedure reproduces the central CT14 Hessian prediction (coinciding with the mean of the MC replicas by construction), and the 68% probability intervals of the MC replicas approximate both the Hessian symmetric and asymmetric uncertainties.

2.3 Generation of Monte-Carlo replicas

We will now present the formulas to generate the Monte-Carlo PDF replicas. We produce N_{rep} sets of the PDF parameters, $\vec{R}^{(k)} \equiv \{R_1^{(k)}, \dots, R_D^{(k)}\}$, which are randomly distributed according to the probability density $\mathcal{P}(\vec{R})$. We will again rely on the experience that $\mathcal{P}(\vec{R})$ is typically close to the standard normal distribution, as in eq. (2.2), with the mean (or, central) values $\vec{R} = \vec{0}$ corresponding to the best-fit PDF set, and the standard deviations that are found from the Hessian PDFs. Therefore the replicas $\vec{R}^{(k)}$ will be sampled from the standard normal distribution in eq. (2.2).

The Monte-Carlo replica $X^{(k)}$ of a QCD observable, X , will be constructed as

$$X^{(k)} = X_0 + d^{(k)} - \Delta, \tag{2.19}$$

in terms of the value X_0 for the central Hessian set, a random shift $d^{(k)} \equiv d(\vec{R}^{(k)})$, and a constant shift Δ applied to all replicas ($k = 1, \dots, N_{\text{rep}}$). Examples of $X(\vec{R})$ include the

PDF, parton luminosity, cross section, or the logarithm of PDF discussed below. The Δ parameter can be set to zero, as in the original Watt-Thorne method. But as was argued in section 2.2, we find it helpful to allow a shift of all replica values by a constant amount Δ .

When X is a PDF, we expect that $\langle X(\vec{R}) \rangle \approx X(\langle \vec{R} \rangle) = X_0$ within the accuracy of the Hessian approximation. [This relation does not hold for an arbitrary X .] Monte-Carlo sampling of the Taylor series may disagree with this condition. The Taylor expansion may lead to systematic biases, as discussed above. That is, X_0 may differ from the mean $\langle X \rangle$ over the replica ensemble,

$$\langle X \rangle \equiv \frac{1}{N_{\text{rep}}} \sum_{k=1}^{N_{\text{rep}}} X^{(k)}, \quad (2.20)$$

even when N_{rep} is large. The mean $\langle X \rangle$ also fluctuates around X_0 , as the number of replicas changes, but these fluctuations are small when N_{rep} is above a few hundred. If our goal is to reproduce the Hessian central value (“truth”) and Hessian uncertainties as closely as possible, we may correct such small offsets of the mean, after all $d^{(k)}$ are computed, by shifting all replicas by a constant amount

$$\Delta \equiv \langle d \rangle = \frac{1}{N_{\text{rep}}} \sum_{k=1}^{N_{\text{rep}}} d^{(k)}. \quad (2.21)$$

We then get $\langle X \rangle = X_0$, up to a small numerical uncertainty, when the replicas are generated according to eqs. (2.19) and (2.21). This choice will be adopted as the default. The random displacements $d^{(k)}$ are found using eq. (2.3), where $R_i = R_i^{(k)}$ for the k -th replica. For the PDFs, $X(\vec{R}) = f_a(x, Q_0; \vec{R})$, this prescription preserves the usual sum rules obeyed by the PDFs, since each replica is a linear combination of the Hessian eigenvector sets, which satisfy the sum rules individually.

Symmetric MC replicas. If only the first derivatives are retained in the Taylor series, then we have

$$d^{(k)} = \sum_{i=1}^D \frac{X_{+i} - X_{-i}}{2} R_i^{(k)}. \quad (2.22)$$

On average, $X^{(k)}$ are symmetrically distributed in the positive and negative directions. The mean of the replicas, remains close to the central Hessian value X_0 : $\langle X \rangle = X_0$ with good accuracy, and the global shift Δ can be neglected. The MC uncertainty on X is then quantified by a standard deviation,

$$\delta^{\text{MC}} X = \sqrt{\langle (X - \langle X \rangle)^2 \rangle}. \quad (2.23)$$

It is expected to approach the 68% c.l. symmetric Hessian uncertainty $\delta_{68}^H X$ in eq. (2.11), in the $N_{\text{rep}} \rightarrow \infty$ limit.

Asymmetric MC replicas. In analogy to eq. (2.12) for the Hessian uncertainties, the *diagonal* second derivatives $\partial^2 X / \partial R_i^2$ can be included as

$$d^{(k)} = \sum_{i=1}^D \left(\frac{X_{+i} - X_{-i}}{2} R_i^{(k)} + \frac{X_{+i} + X_{-i} - 2X_0}{2} \left(R_i^{(k)} \right)^2 \right). \quad (2.24)$$

In this case eq. (2.19) must include the shift term $\Delta = \langle d \rangle$ to satisfy $\langle X \rangle = X_0$. Now the *asymmetric* error estimators are given by the standard deviations that include only positive (negative) displacements from the mean value:

$$\delta^{\text{MC}>X} = +\sqrt{\langle (X - \langle X \rangle)^2 \rangle_{X > \langle X \rangle}}, \quad (2.25)$$

$$\delta^{\text{MC}<X} = -\sqrt{\langle (X - \langle X \rangle)^2 \rangle_{X < \langle X \rangle}}. \quad (2.26)$$

These are the MC counterparts of the asymmetric Hessian uncertainties, $\delta_{68}^{H\gtrless}$, provided by eqs. (2.13) and (2.14).

Compared to the previous work on the MC replicas, our prescription thus includes several new features: the generation of MC replicas using asymmetric displacements (2.24) derived from the Taylor series, a constant shift Δ to ensure agreement between the central Hessian set and the mean of the replicas, and asymmetric standard deviations $\delta^{\text{MC}\gtrless X}$. Watt and Thorne [18] provided the same formulas (2.22), (2.23) for generating the symmetric MC replicas. Their prescription for generating MC replicas with asymmetric displacements is different from the Taylor-series displacements in (2.24) and is given by

$$d^{(k)} = \sum_{i=1}^D \left((X_{+i} - X_0) \Theta [R_i^{(k)} > 0] + (X_{-i} - X_0) \Theta [R_i^{(k)} < 0] \right) |R_i^{(k)}|, \quad (2.27)$$

with $\Delta = 0$.

We advocate using the asymmetric standard deviations (2.25) and (2.26) as estimators for 68% c.l. PDF uncertainties, as they are simple, numerically close to the Hessian asymmetric estimators, cf. the next section, and converge rapidly with the number of replicas.² In terms of the relative significance, using $\delta^{\text{MC}\leq X}$ with Δ shifts, together with the asymmetric standard deviations, is most important for reproducing the asymmetric Hessian uncertainties. As a secondary effect, mild numerical differences were also noticed between the Taylor-series formula (2.24) and Watt-Thorne formula (2.27) for displacements $d^{(k)}$.

2.4 Monte-Carlo replicas for CT14 parton distributions

The formulas derived in section 2.3 are applicable for generating MC replicas for the PDFs themselves, with an extra modification. We perform a truncated Taylor series expansion of $f_a(x, Q; \vec{R})$ at every x point, assuming small deviations from the central value X_0 . In unconstrained x regions, the linear expansion is not sufficient — in fact, it may lead to unphysical solutions such as negative cross sections. Parametrizations for Hessian distributions from the CT family satisfy positivity constraints, $f_a(x, Q) \geq 0$, and render non-negative physical cross sections. To realize the positivity of each MC replica, and eventually to better approximate the non-linear probability distribution, we construct the replicas by Gaussian sampling of $\ln [f_a(x, Q)]$, rather than by sampling $f_a(x, Q)$ directly.

²An alternative estimator, given by the 68% probability interval centered on the mean value, requires sorting the $X^{(k)}$ values and is more susceptible to random fluctuations. Another estimator, given by the maximal displacements of X in the 68% probability hypersphere for \vec{R} , has slow convergence.

The final distribution of CT14 MC replicas includes two families at NNLO in the QCD coupling strength, with 1000 replica sets each, as well as two counterpart families at NLO. Replicas $f_a^{(k)}(x, Q) \equiv f^{(k)}$ of the first type (linear MC, or MC1) are generated assuming normally distributed linear displacements of PDFs, corresponding to eqs. (2.19), (2.24)–(2.26), with $X = f$. Replicas of this type are very similar in spirit to Watt-Thorne replicas; in particular, they may violate positivity.

Replicas of the second type (log MC, or MC2) are generated by sampling of a Gaussian distribution for $X = \ln[f]$, whose variance is estimated as $(\delta^{\text{MC}} X)^2$, and the mean value is set to $\mu = \ln[f_0] - \Delta - (\delta^{\text{MC}} X)^2/2$, where $\Delta = \langle d \rangle$, as above. For the PDFs $f = \exp[X]$, we obtain a log-normal distribution with $\langle f \rangle \approx f_0$. When the PDF error is small, the variance of the PDFs, $\text{var}[f] = f_0^2 \left[e^{(\delta^{\text{MC}} X)^2} - 1 \right]$, becomes equal to the squared Hessian error, $\text{var}[f] \approx (\delta_{68}^H f)^2$. Recall that, if X obeys a normal distribution with the mean μ and variance σ , then the mean and variance of $f = \exp[X]$ are $\langle f \rangle = \exp[\mu + \sigma^2/2]$ and $\text{var}[f] = e^{\mu+2\sigma^2}(e^{\sigma^2} - 1)$. Consequently the “log MC” replicas reproduce the Hessian PDF uncertainty in well-constrained x regions and stay non-negative in the poorly constrained regions.

The difference between the two types of replicas is illustrated by figure 2, showing PDFs for the CT14 NNLO 56 Hessian eigenvector sets (blue circles) and 1000 MC replicas obtained with linear sampling (orange triangles) and logarithmic sampling (red diamonds). The MC replicas are shifted as described above to have $\langle f \rangle_{\text{MC}}$ equal to f_0 , the central PDF of the Hessian set. The vertical axis is scaled as $|f/f_0|^{0.2} \text{sign}(f)$ to better visualize relative variations of both signs in an extended magnitude range.

The value of 1 corresponds to the replica PDF coinciding with that of the central set. The majority of replicas produce PDF values that are close to the central PDF set, within the intervals enclosing 1 where many replica symbols overlap. Some replicas for less constrained PDFs produce very large deviations, corresponding to the values far from 1. The Hessian and log-sampled replicas are non-negative by construction, facilitating positivity of physical cross sections. The linearly sampled replicas can take negative values in extrapolation regions (i.e., lie below the horizontal line at zero). In practice, it is desirable to have non-negative cross sections for individual replicas, even though the uncertainty band does extend to zero at some confidence level.

3 Comparisons of replica sets and Hessian error PDFs

3.1 Estimation of PDF uncertainties

To justify the use of the replica sets for estimating the PDF uncertainties, we must first show that the mean and standard deviation of the replica PDFs themselves are consistent with the results obtained directly from the Hessian error PDFs. A complete set of the comparisons is available at [34]. The PDFs are generated by sampling LHAPDF6 grid files using the MCGEN code [35].

Figure 3 shows the values and error bands, for the ratios of $g(x, Q_0)$, $d(x, Q_0)$, $\bar{u}(x, Q_0)$, and $\bar{s}(x, Q_0)$ PDFs to the respective central CT14 PDFs, at the initial scale $Q_0 = 1.3 \text{ GeV}$,

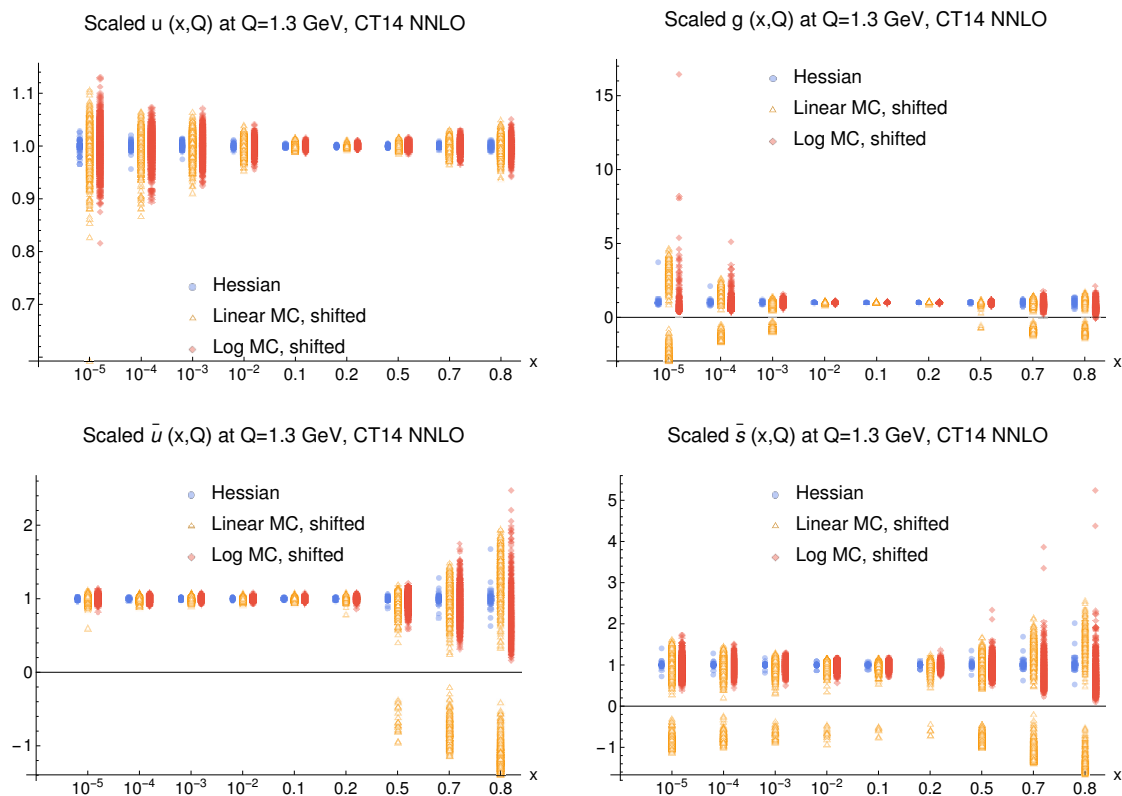


Figure 2. Distributions of individual replicas for MC1 (linear MC, shifted) and MC2 (log MC, shifted) ensembles.

at NNLO. The three green solid curves represent the central PDF (for which the ratio is 1), and the asymmetric upward and downward 68% c.l. deviations calculated from the Hessian master formulae (2.13) and (2.14). The red short-dashed and blue long-dashed curves show the mean and the error bands corresponding to asymmetric standard deviations of the CT14 MC1 and MC2 ensembles with 1000 replicas. We see in the figure that the Hessian and replica PDFs are very consistent in terms of both the means and standard values.³ Both MC1 and MC2 ensembles, obtained by linear and logarithmic sampling of the replicas, result in agreement of comparable quality with the Hessian PDFs at intermediate x . In the extreme x regions, the MC2 error bands tend to be slightly shifted toward positive values. The differences may be somewhat more noticeable for the poorly constrained strangeness PDF, but in all cases the Hessian and MC error bands are fully compatible within the uncertainties of the definition of PDF errors.

In the case of symmetric Hessian uncertainties (2.11) and symmetric standard deviations (2.23), the agreement of Hessian, MC1, and MC2 NNLO ensembles is even better. The corresponding comparisons are shown in figure 4. Equivalent comparisons for CT14 NLO asymmetric errors are presented in figure 5. These figures compare the parton distribution functions at the initial momentum scale Q_0 . If we evolve these to any $Q > Q_0$ using

³The central gluon PDF from the replica set is slightly lower than the Hessian one in the upper left figure 3 at $x < 10^{-3}$ and $Q_0 = 1.3 \text{ GeV}$, where the CT14 gluon vanishes, as a result of numerical roundoff errors.

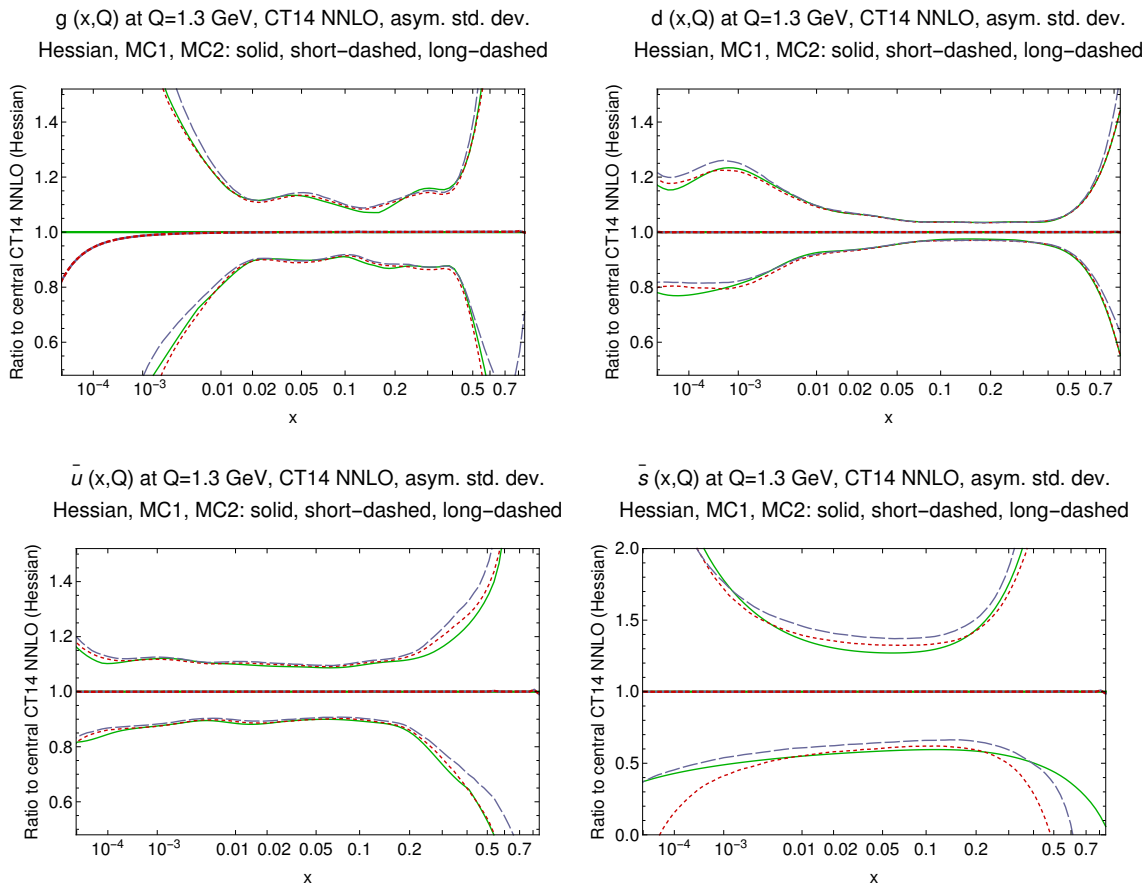


Figure 3. The mean values and asymmetric standard deviations (2.25), (2.26) of the CT14 NNLO MC1 (short-dashed) and MC2 (long-dashed) PDFs, compared to the mean and 68% c.l. uncertainty (eq. (2.11), solid) of the CT14 NNLO Hessian PDF. The PDFs are shown as ratios to the central CT14 fit. Upper panel: $g(x, Q_0)$ and $d(x, Q_0)$. Lower panel: $\bar{u}(x, Q_0)$ and $\bar{s}(x, Q_0)$.

the DGLAP equations, the agreement between the three ensembles further improves. Figures with comparisons of the CT14 Hessian, MC1, and MC2 PDFs at $Q = 2$ and 100 GeV can be viewed at [34].

Lastly, in figure 6, we compare the CT14 NNLO Hessian PDFs to a MC1 version that does not shift the replicas by a constant amount Δ . In this case, the whole bands are shifted compared to the Hessian band across all x , by following their wiggly mean PDFs. [The reader can generate the unshifted replica ensemble with the MCGEN program if desired.] As was argued in the previous section, the Δ shift is introduced to eliminate such spurious variations introduced by truncation of the Taylor series.

It is also useful to compare the *luminosity functions* and their uncertainties, for the MC replicas and the Hessian PDFs. The luminosity function of a parton-parton pair, for production of a final state M^2 at collider energy \sqrt{s} , is defined as in [36], with an additional constraint that the rapidity of the final state, given by $y = \frac{1}{2} \ln(x_2/x_1)$, does not exceed

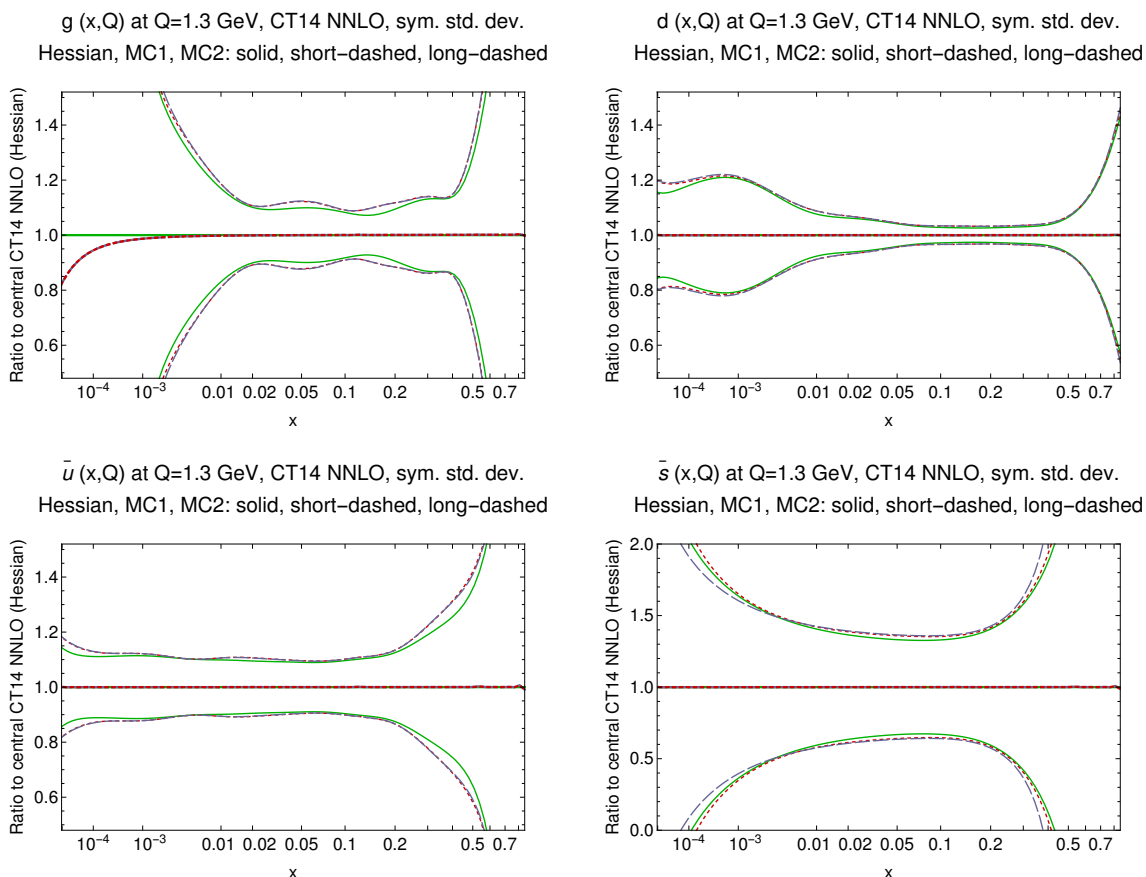


Figure 4. Same as figure 3, for symmetric Hessian uncertainties (2.11) and symmetric MC standard deviations (2.23).

y_{cut} in magnitude:

$$L_{ab}(s, M^2, y_{\text{cut}}) = \frac{1}{1 + \delta_{ab}} \left[\int_{\frac{M}{\sqrt{s}} e^{-y_{\text{cut}}}}^{\frac{M}{\sqrt{s}} e^{y_{\text{cut}}}} \frac{d\xi}{\xi} f_a(\xi, M) f_b\left(\frac{M}{\xi\sqrt{s}}, M\right) + (a \leftrightarrow b) \right]. \quad (3.1)$$

The gluon-gluon and quark-antiquark luminosity functions, calculated from the CT14 Hessian and MC PDFs according to this formula, are shown in figure 7 and 8 as functions of M^2 with $\sqrt{s} = 13$ TeV. We impose a constraint $|y| \leq y_{\text{cut}} = 5$ to exclude contributions to the integral from regions $x < 10^{-5}$, which otherwise may bias the shown luminosities at invariant masses below 40 GeV. For $x < 10^{-5}$, the CT14 PDFs are not constrained by the experimental data, and the final state particles are likely to be produced in the forward region outside of the experimental acceptance of the LHC detectors. With the constraint, the comparison of luminosities is more relevant to the LHC measurements. We see that the Hessian and MC uncertainties agree well across most mass range both in gg and $q\bar{q}$ sectors, with somewhat larger deviations observed at lowest and highest masses.

Taken together, figures 3, 4, and 5 demonstrate that our specifications for the construction of the replica PDFs do yield a satisfactory consistency with the central CT14 PDFs, and with the Hessian error PDFs. Therefore we expect that the uncertainties of

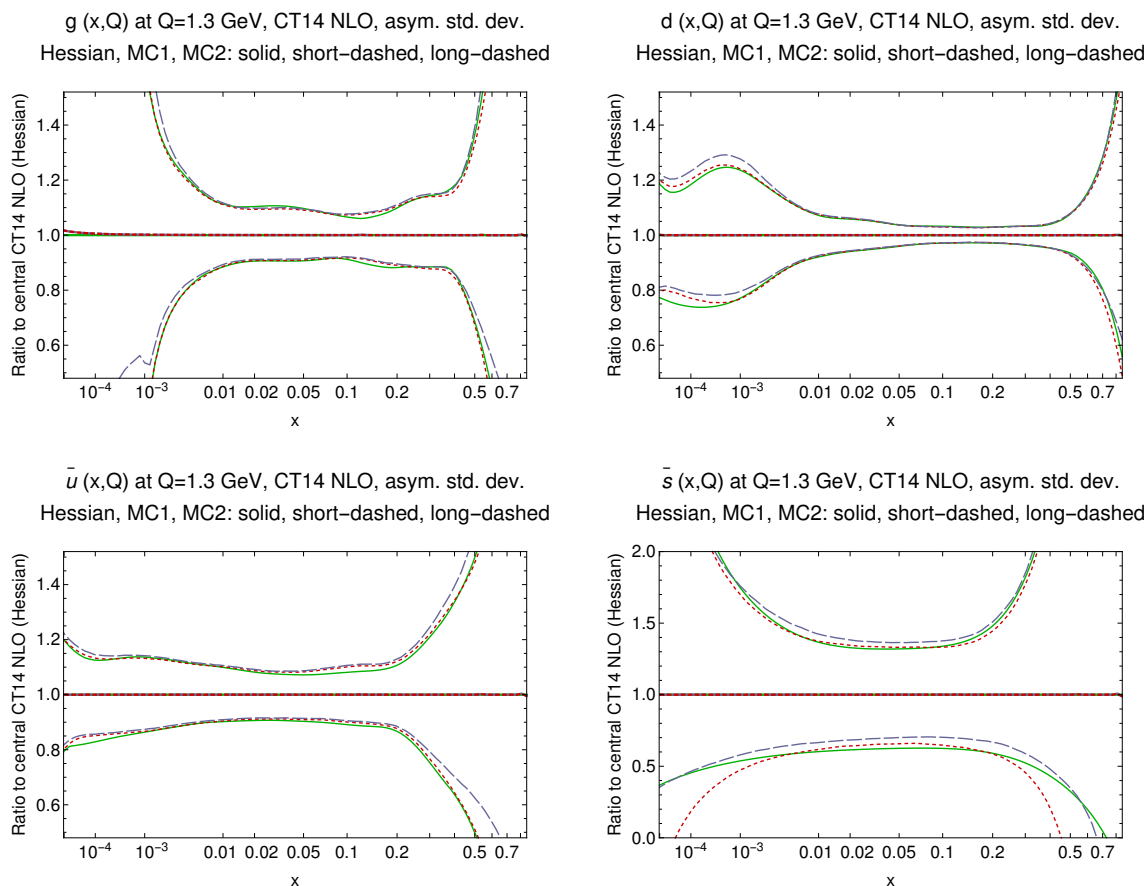


Figure 5. Same as figure 3, at NLO.

cross sections calculated using the replica PDFs will agree with those based on error propagation *via* the Hessian PDFs. Comparisons of the most common LHC cross sections and cross section asymmetries, computed with the CT14 Hessian and MC PDFs, are presented at [34]. By our design, in the regions where the PDFs are well constrained, the ensemble average of the predictions from either CT14MC1 or MC2 PDF set is expected to be close to the central prediction given by the central CT14 Hessian PDF set. Similarly, the asymmetric standard deviations given by the CT14MC PDF set are expected to be very close to the 68% c.l. uncertainties of the CT14 (Hessian) PDF error set, except when the x value is very large (more than 0.7) or small (less than 10^{-5}).

Some examples of cross sections are computed using the APPLGRID fast interface [37] for interpolation of NLO cross sections computed with the help of MADGRAPH_AMC@NLO [38], and AMCFast [39]. The APPLGRID input cross section tables are available at [40] from a study of cross sections based on PDF4LHC15 PDFs [25, 41]. Specifically, we computed CT14 Hessian, MC1 and MC2 rapidity distributions with no or minimal experimental cuts for production W^\pm , Z^0 , $t\bar{t}$, $t\bar{t}\gamma\gamma$, $W^+\bar{c}(W^-c)$ production at the LHC 7,8, and 13 TeV. The renormalization and factorization scales are $\mu_R = \mu_F = M_W$, M_Z , $H_T/2$, $H_T/2$, M_W , respectively. H_T is the scalar sum of transverse masses $\sqrt{p_T^2 + m^2}$

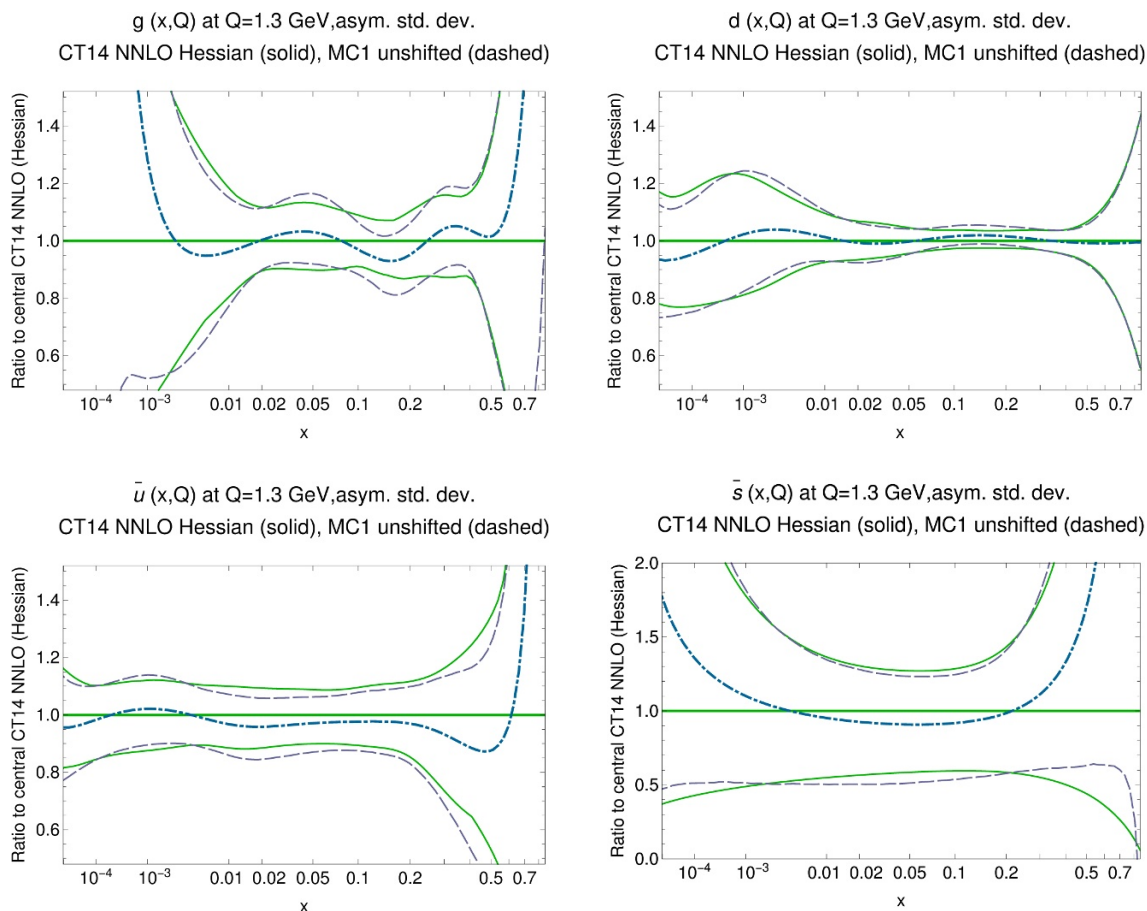


Figure 6. Same as figure 3, for unshifted MC1 and Hessian PDFs of the CT14 NNLO family.

of final-state particles. For $W^+\bar{c}(W^-c)$ production, we neglect small contributions with initial-state c or b quarks. For NLO single-inclusive jet and dijet production, we use public APPLGRID files in the bins of ATLAS measurements [42], created with the program NLO-JET++ [43, 44]. Similarly, the W charge asymmetry in CMS experimental bins [45, 46] is computed with APPLGRID from [47]. An example of the cross sections on the website is shown in figure 9. For ease of comparisons, the PDF uncertainties are plotted both for the cross section values and for ratios to the central CT14 prediction.

3.2 Only an ensemble of MC replicas is meaningful

The results presented above provide a practical prescription for estimating asymmetric Hessian PDF uncertainties using an ensemble of Monte-Carlo replicas. As was already emphasized, some information about the primordial probability distribution in the global fit is lost in the process; however, the prescription is simple and provides a reasonable estimate in most cases. The method is most reliable and unique when the PDF uncertainty is small, meaning that the Taylor series converges fast. When the uncertainty is large, we see more variations, and several alternatives are conceivable. For example, our MC replicas are constructed so that their mean set coincides with the central set of the Hessian ensemble

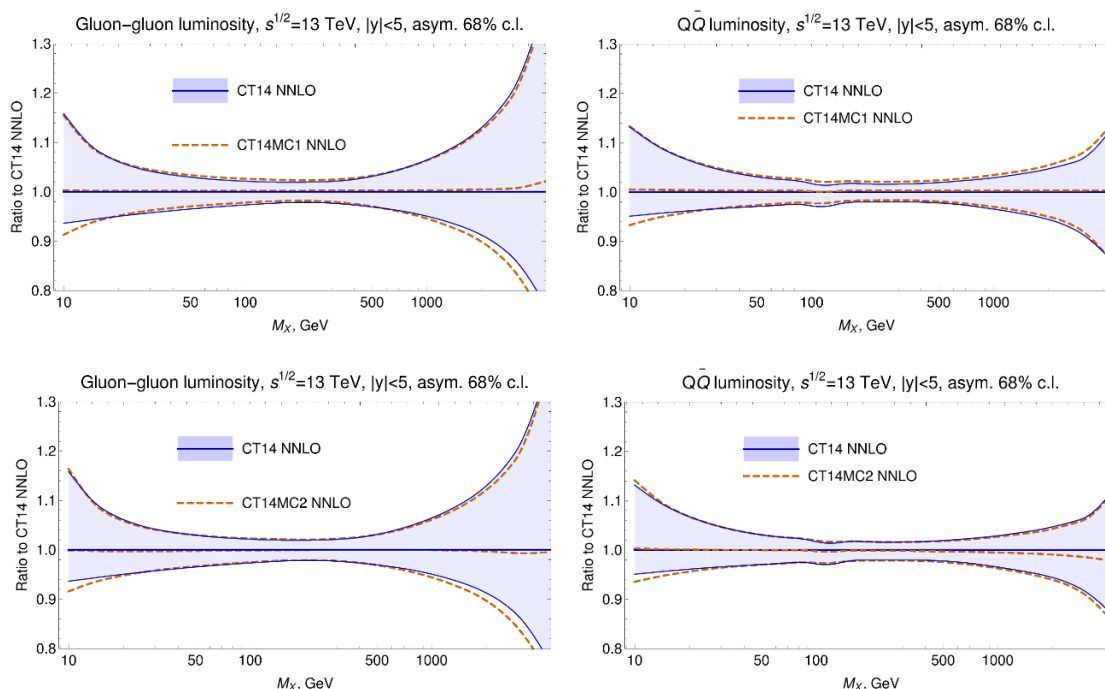


Figure 7. gg and $q\bar{q}$ parton luminosities for the CT14 Hessian (solid blue) and MC (brown dashed) PDF families at NNLO, shown as the ratio to the central Hessian PDF.

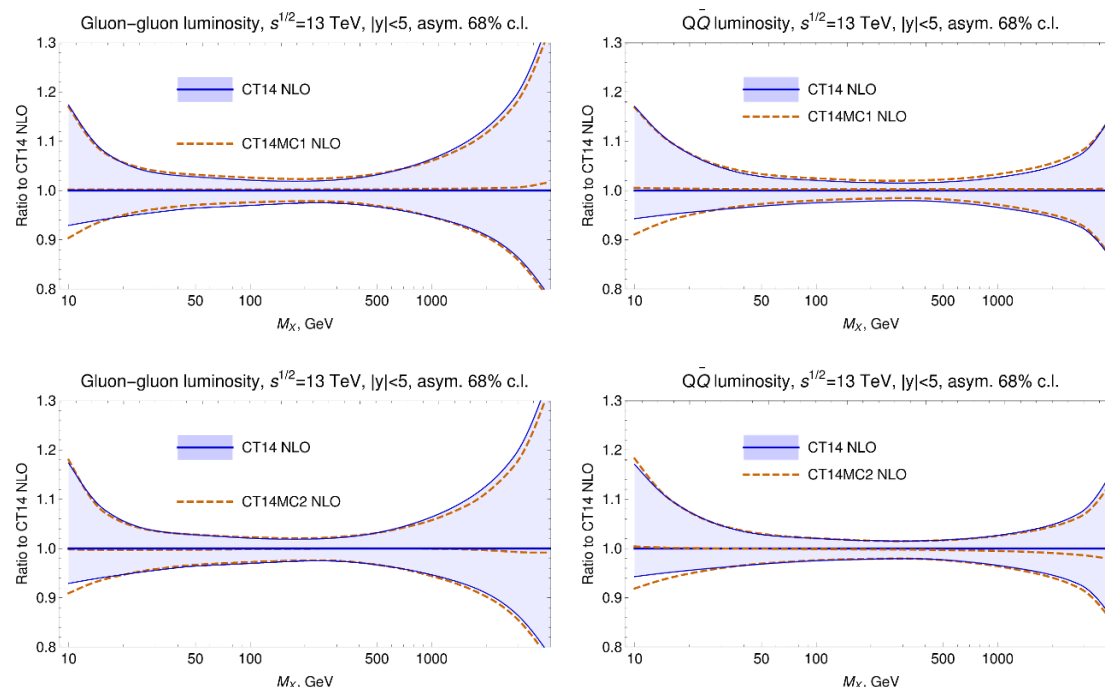


Figure 8. Same as figure 5, for NLO luminosities.




Figure 9. NLO predictions for $d\sigma/dy(W^+)$ in the process $pp \rightarrow W^+ c X$ at the LHC 13 TeV, computed with APPLGRID.

within numerical roundoff errors, as long as x is not too large or too small. Instead, we could choose the central set of the Hessian ensemble to coincide with the median or mode sets of the MC replicas. Similarly, the PDF uncertainty bands may be defined by the positive and negative standard deviations discussed above, or by the boundaries of the 68% c.l. interval centered on the “central PDF set” chosen above. We have checked that the various prescriptions agree well at intermediate x .

It is important to note that most replicas are poor fits to the hadronic data used in the global analysis; however, their averages and standard deviations defined in section 2 provide excellent approximations for the Hessian central PDF set and 68% c.l. uncertainties. This is demonstrated in figure 10, showing histograms of χ^2 values for the global data (3174 data points; left panel) and for the combined HERA-1 data (579 data points [48]; right panel), for the 1000 replicas in the CT14 NNLO MC1 and MC2 ensembles. The vast majority of replicas yield very large χ^2 values for the global data, and even for the single experiment. The random fluctuations of the individual replicas, which result in large χ^2 values for any single replica, will largely cancel in the ensemble averages.

It is straightforward to understand why the result shown in figure 10 occurs. Imagine that we construct a D -dimensional vector \vec{a} whose coordinates are given by random variates a_i sampled from a standard normal distribution. The length of \vec{a} will often turn out to be significantly larger than 1. If those parameter values are used for the PDFs, $\chi^2(\vec{a})$ will tend to be much larger than the minimal χ^2 in the fit, especially if the number of dimensions D is large. [Recall that the volume of a D -dimensional unit ball vanishes in the limit of large D . The vast majority of replicas will have several parameters far outside of the unit ball, i.e., away from the best fit by many standard deviations.]

The expected value of the length of \vec{a} can be easily found as

$$\begin{aligned}
 \langle |\vec{a}| \rangle &= \frac{\int_{-\infty}^{\infty} da_1 \int_{-\infty}^{\infty} da_2 \dots \int_{-\infty}^{\infty} da_D \sqrt{\sum_{i=1}^D a_i^2} \exp\left(-\sum_{i=1}^D a_i^2/2\right)}{\int_{-\infty}^{\infty} da_1 \int_{-\infty}^{\infty} da_2 \dots \int_{-\infty}^{\infty} da_D \exp\left(-\sum_{i=1}^D a_i^2/2\right)} \\
 &= (2\pi)^{-D/2} \cdot \int_0^{\infty} da a^D \exp(-a^2/2) \cdot \int d\Omega_D \\
 &= \frac{\sqrt{2}\Gamma((D+1)/2)}{\Gamma(D/2)} \approx \sqrt{D}.
 \end{aligned}
 \tag{3.2}$$

The final approximation in (3.2) follows from Stirling’s formula for the gamma function. For $D = 28$, eq. (3.2) gives $\langle |\vec{a}| \rangle \approx 5.24$; that is, a typical displacement vector of a CT14 replica is more than five standard units in length. As a result, most replicas have a χ^2 value that significantly exceeds the CT14 best-fit value of about 3250 for the global data and 590 for the HERA-1 data. If the standard deviation of the normal distribution is rather about 6 units (corresponding to $\Delta\chi^2 \approx 100$ at 90% c.l.), the average displacement vector corresponds to $\Delta\chi^2 \approx (5.24 \cdot 6)^2 \approx 1000$ for the global data and $\Delta\chi^2 \approx 180$ (scaled down by 579/3174) for the HERA-1 data. This is far outside the typical 1σ for the χ^2 distribution, given by $T^2 \approx 36$, $\sqrt{2N_{pts}} \approx 80(35)$, or another common estimator! These estimates are in good agreement with the actual χ^2 averages over the CT14 replicas, denoted as $\langle \chi^2 \rangle_{\text{rep}}$, and equal in the case of the MC1 (MC2) NNLO ensembles to ≈ 4300 (4200) for the global χ^2 , and ≈ 720 (730) for the HERA-1 χ^2 .

It is interesting to note [49] that the χ^2 distribution for the CT14 MC replicas is quite similar in shape to the χ^2 distributions for the global data of NNPDF replicas, obtained with an entirely different methodology [3, 14–17]. For example, for the four NNPDF3.0 NLO fits listed in table 1 of ref. [50], the equivalent of our $\langle \chi^2 \rangle_{\text{rep}}$ is 600–1000 units higher than the χ^2 value for the average PDF set of all replicas.⁴ This can be qualitatively understood by noticing that the tensions between the individual experiments can be effectively accommodated by introducing a tolerance $T_{\text{exp}} \sim 2$ on the global χ^2 at the 68% c.l. [26, 51, Section 7], and that the neural network parametrizations effectively have of order $D = 100 - 250$ free parameters. This predicts $\langle \chi^2 \rangle_{\text{rep}} - \chi_{\text{min}}^2 \sim T_{\text{exp}}^2 D = 400 - 1000$, in fair agreement with the actual NNPDF3.0 outcomes.

Individual replica PDFs thus cannot be considered as alternative PDFs with approximately the same accuracy as the central fit, or even with the accuracy of the Hessian error PDFs. The replica PDFs are meaningful only inside an ensemble, predictions based on them must be calculated by averaging over the ensemble.

4 Conclusions

The paper explored methods for the conversion of Hessian PDF error sets into Monte-Carlo replica PDFs. We observed that parameters of the PDFs approximately obey a multi-dimensional normal distribution in the vicinity of the global χ^2 minimum. A statistical sample of MC replicas dependent on the probability distribution in the D -dimensional parameter space can be reconstructed under the Gaussian assumption from the distribution of the Hessian error PDFs on the $(D - 1)$ -dimensional boundary of some confidence region enclosing the best fit. The CT14 Hessian error PDFs allow the user to estimate the PDF errors and account for the asymmetries due to mild deviations from the Gaussian shape and nonlinear nature of PDF parametrization. The CT14 PDFs are parameterized to be explicitly non-negative, as desired for obtaining sensible cross sections and uncertainties. In section 2 we have generalized the algorithms for the generation of MC replicas and

⁴The NNPDF table constructs $\langle \chi^2 \rangle_{\text{rep}}$ from the χ^2 values between theoretical predictions for an individual PDF replica and true (not fluctuated) experimental data. The average is over the ensemble of the NNPDF replicas.

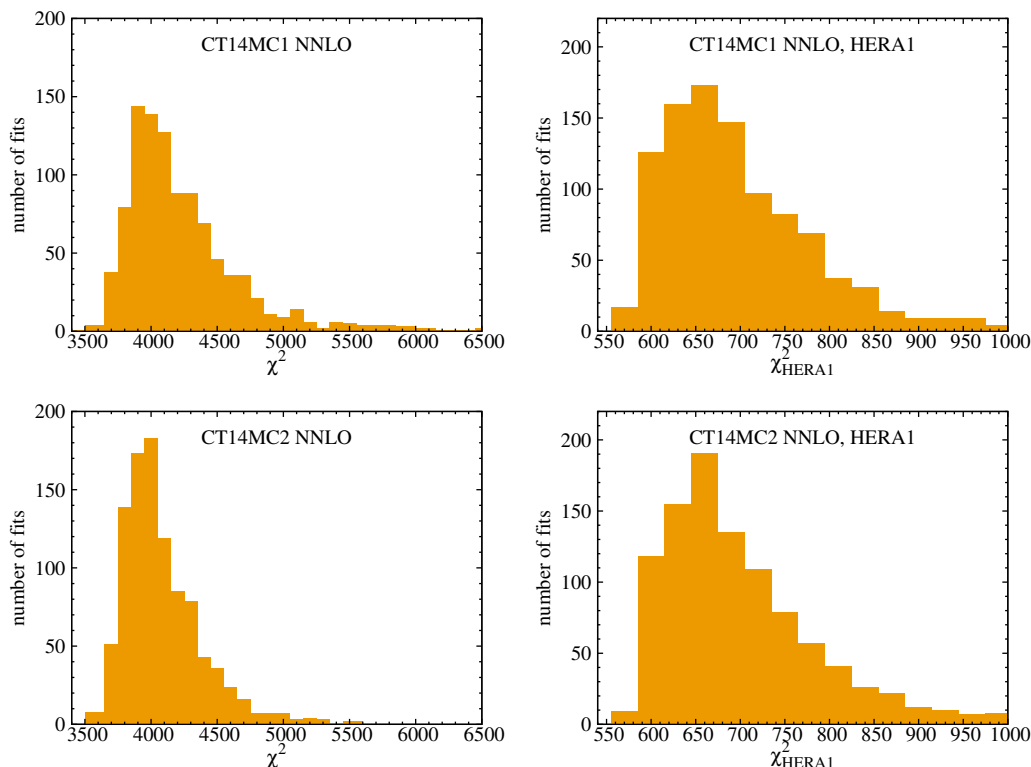


Figure 10. χ^2 for global data (left) and combined HERA-1 data (right), for 1000 replicas of the CT14MC NNLO ensemble.

for the estimation of PDF uncertainties necessary to reproduce these features. Based on these results, we have generated two replica ensembles, designated as CT14MC1 and CT14MC2 PDFs. The average PDF sets and the asymmetric standard deviations of MC1 and MC2 closely reproduce the central sets and the 68% c.l. asymmetric uncertainties of the respective Hessian ensembles, cf. section 3. With the naive conversion, the mean set over the replicas may deviate from the central Hessian PDF set because of the truncation of Taylor series. After we apply a correction, the MC mean and central Hessian sets coincide, preventing potential discrepancies in the predictions. In addition, the CT14MC2 PDF replicas, obtained by sampling of a log-normal distribution, are explicitly non-negative.

To achieve good agreement between the CT14 Hessian and MC uncertainty bands, we recommend to construct the MC replicas by random sampling according to the algorithm summarized as eqs. (2.19), (2.21), (2.24), and section 2.4. The central PDF set of such MC ensemble by construction coincides with the central Hessian PDF set. The PDF uncertainties for the MC ensemble are to be estimated by asymmetric standard deviations (2.25), (2.26), introduced as the counterparts for the master formulas (2.13), (2.14) for the Hessian asymmetric errors.

Our results in section 2 clarify several aspects of the replica generation method that were not addressed in the previous work. Section 3 demonstrates that, in most of the x range, where the PDF uncertainties are small, we observe good agreement between the

CT14 Hessian, MC1, and MC2 error bands. In the extrapolation regions, where the linear approximations cease to be adequate, differences between the Hessian, MC1 and MC2 error bands are indicative of intrinsic ambiguities in the replica generation methods. At such extreme x , they may affect the combination of the Hessian and MC ensembles, as in the recommended PDF4LHC procedure, or PDF replica re-weighting. These ambiguities require consideration when the converted replicas are utilized in PDF reweighting or PDF combinations. The positivity constraint on CT14 MC2 yields more physical behavior in poorly constrained x intervals.

Once the MC replicas are generated by conversion, we examine their statistical properties in section 3.2. We point out that many individual MC replica sets yield poor χ^2 , so only their statistical combinations, such as the mean, standard deviation, etc., are meaningful in applications. [The majority of MC replicas aren't good fits, their combination is.] On the other hand, the χ^2 distributions of MC replicas obtained by conversion (such as CT14 MC) or genetic algorithm (such as NNPDF) are very similar. This reflects the underlying commonalities of the two methods, leading to comparable PDF errors in the two approaches in spite of the distinct error definitions and fitting procedures.

The CT14 MC1 and MC2 ensembles at NNLO and NLO accuracy, together with a fast standalone driver program for their interpolation, can be downloaded at [52]. They are also distributed as a component of the LHAPDF6 library [53]. A public C++ code MCGEN is made available [35] for generation of MC PDF replicas using the normal, log-normal, and Watt-Thorne sampling methods, and with or without including the Δ shifts. MCGEN can be run as a standalone program or together with the Mathematica package MP4LHC for combination of PDF ensembles according to the meta-parametrization method [19]. After N_{input} Hessian error PDF sets are read in the form of LHAPDF6 grids, N_{rep} output replicas are generated by random displacements of the Hessian replicas. Besides the replica generation, MCGEN supports various algebraic operations with PDFs in the format of LHAPDF6 grid files, such as addition, averaging, and multiplication of the tables in which the PDF values are stored.

Acknowledgments

We thank R. Thorne for an instructive discussion about statistical properties of MC replicas. J.H., J.G., and P.N. thank the Kavli Institute for Theoretical Physics at Santa Barbara, CA and organizers of the ‘‘LHC Run II and Precision Frontier’’ research program for hospitality and productive atmosphere during completion of this paper. This research was supported in part by the National Science Foundation under Grants No. PHY-1410972 and PHY11-25915; by the U.S. Department of Energy under Contract No. DE-AC02-06CH11357 and Grants DE-SC0013681 and DE-SC0010129; and by the National Natural Science Foundation of China under Grant No. 11465018.

A Definition of χ^2 in CT fits

For completeness of the presentation, in this appendix, we reproduce the definition for the figure-of-merit function χ^2 and the procedure for determination of PDF uncertainties in the CT14 fit, introduced in full detail in refs. [1, 26, 27].

The most probable solutions for CT14 PDFs are found by minimization of a *global* log-likelihood function

$$\chi_{\text{global}}^2 = \sum_{n=1}^{N_{\text{exp}}} \chi_n^2 + \chi_{th}^2, \quad (\text{A.1})$$

which sums contributions χ_n^2 from N_{exp} fitted experiments, and a contribution χ_{th}^2 specifying theoretical conditions (“Lagrange Multiplier constraints”) imposed on some PDF parameters. In turn, the terms χ_n^2 are given by

$$\chi_n^2(\vec{a}, \vec{\lambda}) = \chi_D^2 + \chi_\lambda^2, \quad (\text{A.2})$$

where

$$\chi_D^2 \equiv \sum_{k=1}^{N_{pts}} \frac{1}{s_k^2} \left(D_k - T_k(\vec{a}) - \sum_{\alpha=1}^{N_\lambda} \beta_{k,\alpha} \lambda_\alpha \right)^2, \quad (\text{A.3})$$

and

$$\chi_\lambda^2 \equiv \sum_{\alpha=1}^{N_\lambda} \lambda_\alpha^2. \quad (\text{A.4})$$

The χ_n^2 contribution is a function of the PDF parameters \vec{a} and systematic nuisance parameters $\vec{\lambda}$. For a k -th data point, T_k , D_k , s_k , and $\beta_{k\alpha}$ are the theoretical prediction, central experimental value, uncorrelated experimental uncertainty, and systematic correlation matrix, respectively.

The minimum of the χ_{global}^2 function is found iteratively by the method of steepest descent using the program MINUIT. The boundaries of the 90% c.l. region around the minimum of χ_{global}^2 , and the eigenvector PDF sets quantifying the associated uncertainty, are found by iterative diagonalization of the Hessian matrix [7, 10], which finds independent, or eigenvector, combinations of the PDF parameters a_i . These combinations are denoted by z_i , with the best-fit parameter combination corresponding to $z_1 = z_2 = \dots = 0$. The 90% c.l. boundary around the best fit is determined by applying two tiers of criteria, based on the increase in the global χ_{global}^2 summed over all experiments, and the agreement with individual experimental data sets [26, 27, 54]. The first type of condition demands that the global χ^2 does not increase above the best-fit value by more than $\Delta\chi^2 = T^2$, where the 90% c.l. region corresponds to $T \approx 10$. The second condition introduces a penalty term P in χ^2 when establishing the confidence region, which quickly grows when the fit ceases to agree with any specific experiment within the 90% c.l. for that experiment. The effective function $\chi_{\text{eff}}^2 = \chi_{\text{global}}^2 + P$ is scanned along each eigenvector direction until χ_{eff}^2 increases above the tolerance bound, or rapid χ_{eff}^2 growth due to the penalty P is triggered.

The penalty term is constructed as

$$P = \sum_{n=1}^{N_{\text{exp}}} (S_n)^p \theta(S_n) \quad (\text{A.5})$$

from the equivalent Gaussian variables S_n that obey an approximate standard normal distribution independently of the number of data points $N_{pts,n}$ in the experiment. Every

S_n is a monotonically increasing function of the respective χ_n^2 given in [54, 55]. The power $p = 16$ is chosen so that $(S_n)^p$ sharply increases from zero when S_n approaches 1.3, the value corresponding to 90% c.l. cutoff.

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

References

- [1] S. Dulat et al., *New parton distribution functions from a global analysis of quantum chromodynamics*, *Phys. Rev. D* **93** (2016) 033006 [[arXiv:1506.07443](https://arxiv.org/abs/1506.07443)] [[INSPIRE](#)].
- [2] L.A. Harland-Lang, A.D. Martin, P. Motylinski and R.S. Thorne, *Parton distributions in the LHC era: MMHT 2014 PDFs*, *Eur. Phys. J. C* **75** (2015) 204 [[arXiv:1412.3989](https://arxiv.org/abs/1412.3989)] [[INSPIRE](#)].
- [3] NNPDF collaboration, R.D. Ball et al., *Parton distributions for the LHC Run II*, *JHEP* **04** (2015) 040 [[arXiv:1410.8849](https://arxiv.org/abs/1410.8849)] [[INSPIRE](#)].
- [4] ZEUS, H1 collaboration, H. Abramowicz et al., *Combination of measurements of inclusive deep inelastic $e^\pm p$ scattering cross sections and QCD analysis of HERA data*, *Eur. Phys. J. C* **75** (2015) 580 [[arXiv:1506.06042](https://arxiv.org/abs/1506.06042)] [[INSPIRE](#)].
- [5] S. Alekhin, J. Blumlein and S. Moch, *The ABM parton distributions tuned to LHC data*, *Phys. Rev. D* **89** (2014) 054028 [[arXiv:1310.3059](https://arxiv.org/abs/1310.3059)] [[INSPIRE](#)].
- [6] A. Accardi, L.T. Brady, W. Melnitchouk, J.F. Owens and N. Sato, *Constraints on large- x parton distributions from new weak boson production and deep-inelastic scattering data*, *Phys. Rev. D* **93** (2016) 114017 [[arXiv:1602.03154](https://arxiv.org/abs/1602.03154)] [[INSPIRE](#)].
- [7] J. Pumplin et al., *Uncertainties of predictions from parton distribution functions. 2. The Hessian method*, *Phys. Rev. D* **65** (2001) 014013 [[hep-ph/0101032](https://arxiv.org/abs/hep-ph/0101032)] [[INSPIRE](#)].
- [8] W.T. Giele and S. Keller, *Implications of hadron collider observables on parton distribution function uncertainties*, *Phys. Rev. D* **58** (1998) 094023 [[hep-ph/9803393](https://arxiv.org/abs/hep-ph/9803393)] [[INSPIRE](#)].
- [9] W.T. Giele, S.A. Keller and D.A. Kosower, *Parton distribution function uncertainties*, [hep-ph/0104052](https://arxiv.org/abs/hep-ph/0104052) [[INSPIRE](#)].
- [10] D. Stump et al., *Uncertainties of predictions from parton distribution functions. 1. The Lagrange multiplier method*, *Phys. Rev. D* **65** (2001) 014012 [[hep-ph/0101051](https://arxiv.org/abs/hep-ph/0101051)] [[INSPIRE](#)].
- [11] A.M. Cooper-Sarkar, *Uncertainties on parton distribution functions from the ZEUS NLO QCD fit to data on deep inelastic scattering*, *J. Phys. G* **28** (2002) 2669 [[hep-ph/0205153](https://arxiv.org/abs/hep-ph/0205153)] [[INSPIRE](#)].
- [12] J. Pumplin, D.R. Stump, J. Huston, H.L. Lai, P.M. Nadolsky and W.K. Tung, *New generation of parton distributions with uncertainties from global QCD analysis*, *JHEP* **07** (2002) 012 [[hep-ph/0201195](https://arxiv.org/abs/hep-ph/0201195)] [[INSPIRE](#)].
- [13] P.M. Nadolsky and Z. Sullivan, *PDF uncertainties in WH production at Tevatron*, *eConf C* **010630** (2001) P510 [[hep-ph/0110378](https://arxiv.org/abs/hep-ph/0110378)] [[INSPIRE](#)].
- [14] S. Forte, L. Garrido, J.I. Latorre and A. Piccione, *Neural network parametrization of deep inelastic structure functions*, *JHEP* **05** (2002) 062 [[hep-ph/0204232](https://arxiv.org/abs/hep-ph/0204232)] [[INSPIRE](#)].
- [15] NNPDF collaboration, L. Del Debbio et al., *Neural network determination of parton distributions: the nonsinglet case*, *JHEP* **03** (2007) 039 [[hep-ph/0701127](https://arxiv.org/abs/hep-ph/0701127)] [[INSPIRE](#)].

- [16] NNPDF collaboration, R.D. Ball et al., *A determination of parton distributions with faithful uncertainty estimation*, *Nucl. Phys. B* **809** (2009) 1 [Erratum *ibid.* **B 816** (2009) 293] [[arXiv:0808.1231](#)] [[INSPIRE](#)].
- [17] R.D. Ball et al., *Parton distributions with LHC data*, *Nucl. Phys. B* **867** (2013) 244 [[arXiv:1207.1303](#)] [[INSPIRE](#)].
- [18] G. Watt and R.S. Thorne, *Study of Monte Carlo approach to experimental uncertainty propagation with MSTW 2008 PDFs*, *JHEP* **08** (2012) 052 [[arXiv:1205.4024](#)] [[INSPIRE](#)].
- [19] J. Gao and P. Nadolsky, *A meta-analysis of parton distribution functions*, *JHEP* **07** (2014) 035 [[arXiv:1401.0013](#)] [[INSPIRE](#)].
- [20] S. Carrazza, S. Forte, Z. Kassabov, J.I. Latorre and J. Rojo, *An unbiased Hessian representation for Monte Carlo PDFs*, *Eur. Phys. J. C* **75** (2015) 369 [[arXiv:1505.06736](#)] [[INSPIRE](#)].
- [21] S. Carrazza, S. Forte, Z. Kassabov and J. Rojo, *Specialized minimal PDFs for optimized LHC calculations*, *Eur. Phys. J. C* **76** (2016) 205 [[arXiv:1602.00005](#)] [[INSPIRE](#)].
- [22] NNPDF collaboration, R.D. Ball et al., *Reweighting NNPDFs: the W lepton asymmetry*, *Nucl. Phys. B* **849** (2011) 112 [Erratum *ibid.* **B 854** (2012) 926] [[arXiv:1012.0836](#)] [[INSPIRE](#)].
- [23] R.D. Ball et al., *Reweighting and unweighting of parton distributions and the LHC W lepton asymmetry data*, *Nucl. Phys. B* **855** (2012) 608 [[arXiv:1108.1758](#)] [[INSPIRE](#)].
- [24] S. Carrazza, J.I. Latorre, J. Rojo and G. Watt, *A compression algorithm for the combination of PDF sets*, *Eur. Phys. J. C* **75** (2015) 474 [[arXiv:1504.06469](#)] [[INSPIRE](#)].
- [25] J. Butterworth et al., *PDF4LHC recommendations for LHC Run II*, *J. Phys. G* **43** (2016) 023001 [[arXiv:1510.03865](#)] [[INSPIRE](#)].
- [26] H.-L. Lai et al., *New parton distributions for collider physics*, *Phys. Rev. D* **82** (2010) 074024 [[arXiv:1007.2241](#)] [[INSPIRE](#)].
- [27] J. Gao et al., *CT10 next-to-next-to-leading order global analysis of QCD*, *Phys. Rev. D* **89** (2014) 033009 [[arXiv:1302.6246](#)] [[INSPIRE](#)].
- [28] A. Belyaev, J. Pumplin, W.-K. Tung and C.P. Yuan, *Uncertainties of the inclusive Higgs production cross section at the Tevatron and the LHC*, *JHEP* **01** (2006) 069 [[hep-ph/0508222](#)] [[INSPIRE](#)].
- [29] S. Dulat et al., *Higgs boson cross section from CTEQ-TEA global analysis*, *Phys. Rev. D* **89** (2014) 113002 [[arXiv:1310.7601](#)] [[INSPIRE](#)].
- [30] C. Schmidt et al., *Uncertainties on Higgs and $t\bar{t}$ predictions at the LHC from CTEQ-TEA Global Analysis*, *PoS(DIS2014)146*.
- [31] J. Pumplin, J. Huston, H.L. Lai, P.M. Nadolsky, W.-K. Tung and C.P. Yuan, *Collider inclusive jet data and the gluon distribution*, *Phys. Rev. D* **80** (2009) 014019 [[arXiv:0904.2424](#)] [[INSPIRE](#)].
- [32] J.F. Kenney and E.S. Keeping, *Mathematics of statistics*, 3rd edition, Van Nostrand, New York U.S.A. (1954), see chapters 4.7-4.9.
- [33] <http://mathworld.wolfram.com/Mode.html>.
- [34] <http://hep.pa.msu.edu/cteq/public/ct14/MC/>.
- [35] <http://metapdf.hepforge.org/mcgen/>.

- [36] J.M. Campbell, J.W. Huston and W.J. Stirling, *Hard interactions of quarks and gluons: a primer for LHC physics*, *Rept. Prog. Phys.* **70** (2007) 89 [[hep-ph/0611148](#)] [[INSPIRE](#)].
- [37] T. Carli et al., *A posteriori inclusion of parton density functions in NLO QCD final-state calculations at hadron colliders: the APPLGRID project*, *Eur. Phys. J. C* **66** (2010) 503 [[arXiv:0911.2985](#)] [[INSPIRE](#)].
- [38] J. Alwall et al., *The automated computation of tree-level and next-to-leading order differential cross sections and their matching to parton shower simulations*, *JHEP* **07** (2014) 079 [[arXiv:1405.0301](#)] [[INSPIRE](#)].
- [39] V. Bertone, R. Frederix, S. Frixione, J. Rojo and M. Sutton, *aMCfast: automation of fast NLO computations for PDF fits*, *JHEP* **08** (2014) 166 [[arXiv:1406.7693](#)] [[INSPIRE](#)].
- [40] <http://metapdf.hepforge.org/2016.pdf4lhc/>.
- [41] J.R. Andersen et al., *Les Houches 2015: Physics at TeV colliders standard model working group report*, [arXiv:1605.04692](#) [[INSPIRE](#)].
- [42] ATLAS collaboration, *Measurement of inclusive jet and dijet production in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector*, *Phys. Rev. D* **86** (2012) 014022 [[arXiv:1112.6297](#)] [[INSPIRE](#)].
- [43] Z. Nagy, *Three jet cross-sections in hadron hadron collisions at next-to-leading order*, *Phys. Rev. Lett.* **88** (2002) 122003 [[hep-ph/0110315](#)] [[INSPIRE](#)].
- [44] Z. Nagy, *Next-to-leading order calculation of three jet observables in hadron hadron collision*, *Phys. Rev. D* **68** (2003) 094002 [[hep-ph/0307268](#)] [[INSPIRE](#)].
- [45] CMS collaboration, *Measurement of the electron charge asymmetry in inclusive W production in pp collisions at $\sqrt{s} = 7$ TeV*, *Phys. Rev. Lett.* **109** (2012) 111806 [[arXiv:1206.2598](#)] [[INSPIRE](#)].
- [46] CMS collaboration, *Measurement of the muon charge asymmetry in inclusive $pp \rightarrow W + X$ production at $\sqrt{s} = 7$ TeV and an improved determination of light parton distribution functions*, *Phys. Rev. D* **90** (2014) 032004 [[arXiv:1312.6283](#)] [[INSPIRE](#)].
- [47] S. Alekhin et al., *HERAFitter*, *Eur. Phys. J. C* **75** (2015) 304 [[arXiv:1410.4412](#)] [[INSPIRE](#)].
- [48] ZEUS, H1 collaboration, F.D. Aaron et al., *Combined measurement and QCD analysis of the inclusive $e^\pm p$ scattering cross sections at HERA*, *JHEP* **01** (2010) 109 [[arXiv:0911.0884](#)] [[INSPIRE](#)].
- [49] R. Thorne, *Summary of structure functions and PDFs working group*, in the summary talk of the *XXIV Workshop on Deep-Inelastic Scattering (DIS) and Related Subjects (DIS'2016)*, Aprile 11–15, Hamburg, Germany (2016).
- [50] NNPDF collaboration, R.D. Ball et al., *A determination of the charm content of the proton*, *Eur. Phys. J. C* **76** (2016) 647 [[arXiv:1605.06515](#)] [[INSPIRE](#)].
- [51] J. Pumplin, *Experimental consistency in parton distribution fitting*, *Phys. Rev. D* **81** (2010) 074010 [[arXiv:0909.0268](#)] [[INSPIRE](#)].
- [52] <http://http://hep.pa.msu.edu/cteq/public/>.
- [53] <http://lhpdf.hepforge.org>.
- [54] S. Dulat et al., *Intrinsic charm parton distribution functions from CTEQ-TEA global analysis*, *Phys. Rev. D* **89** (2014) 073004 [[arXiv:1309.0025](#)] [[INSPIRE](#)].
- [55] T. Lewis, *A simple improved-accuracy normal approximation for χ^2* , *Austral. J. Statist.* **30A** (1988) 160.