We thank the referee for the close reading of our manuscript and for the detailed report which will help us greatly in clarifying various issues. We have addressed the various comments raised by the referee, as we now describe in turn:

1. *Page 3, first full paragraph. The discussion of the approach to calculating the photon PDF based on a theoretically motivated ansatz gives undue prominence to CT14QED. While it is true that the final public release of the CT14QED analysis includes the elastic component of the photon PDF, the original study did not consider this. The inclusion of the elastic component came almost a year after the original release, and came subsequently to the discussion of [33,34]. In addition, Refs [30-34] all include a model of the inelastic component, and so are qualitatively no different in approach from the CT14QED set. Given that the idea of this introductory paragraph is to describe the model dependent approach, CT14QED and Refs [30-34] should be dealt with on a more equal footing, ideally giving some indication of how these ideas have developed chronologically.*

   We agree with the referee on this point. We have restructured the introduction to highlight more properly the original contributions of CT14QED as compared to those of Refs. [30-34].

2. *Page 3, last sentence of second paragraph. English-wise this needs a little rewording: "Although this dataset is particularly...". It would also perhaps be fairer to [39] to say that some reduction in uncertainty is achieved relative to the baseline.*

   We have readjusted the text following the suggestion from the referee.

3. *Page 3. The point should be made somewhere here that the elastic component is by far the dominant contribution to the input photon distribution, in particular at higher x, and thus the uncertainties are already greatly reduced by including this. This is even briefly discussed later on at the end of section 2.4, but is easily missed there.*

   We have readjusted the text following the suggestion from the referee.

4. *Page 3, third full paragraph. Here or somewhere else, the earlier works (Anlauf et al. Comput.Phys.Commun. 70 (1992) 97-119, Mukherjee and Pisano Eur.Phys.J. C30 (2003) 477-486, Blumlein et al J.Phys. G19 (1993) 1695-1703) should be referenced to. These independently calculated expressions using a similar approach to LUXqed, i.e. relating the inelastic photon to the proton structure functions. These resulted in expressions for the photon that were very close to the LUXqed result, with the exception of the limits on the Q^2 integral, which were not correct, and the missing mass correction in the Blumlein case. Clearly LUXqed represents the state of the art in this respect, but a reference and brief description would be fair.*

   We agree with the referee on this point, and we have added to the manuscript the requested suggestions.

5. *Page 3, second to last paragraph. The statement that the new photon is 'fully consistent' is not supported by the current results, even when phrased in terms of the impact of the photon PDF. From Fig. 4.3 (also 3.3) we can see that the 3.1lux photon-initiated contribution is important relative to other PDF uncertainties, but also inconsistent with the earlier 3.0 prediction. I discuss this more below, but this cannot be the right thing to say here if this result stands.*

   We address this point in our answer to point 12) below.

6. *Section 2.3, third paragraph. It would be useful to show a plot of the impact of the higher order corrections on the photon-photon luminosity.*

We agree with the referee on this point, and we have a added to Section 2 a new plot showing the impact of the higher order corrections to the PDF luminosity using NNPDF3.1luxQED. For completeness, we have also added the corresponding plot for the impact of the higher order QED corrections to the deep inelastic structure functions.

The plot that has been added to the paper, which is Fig. 2.1 in the revised version of the manuscript, is the following one:

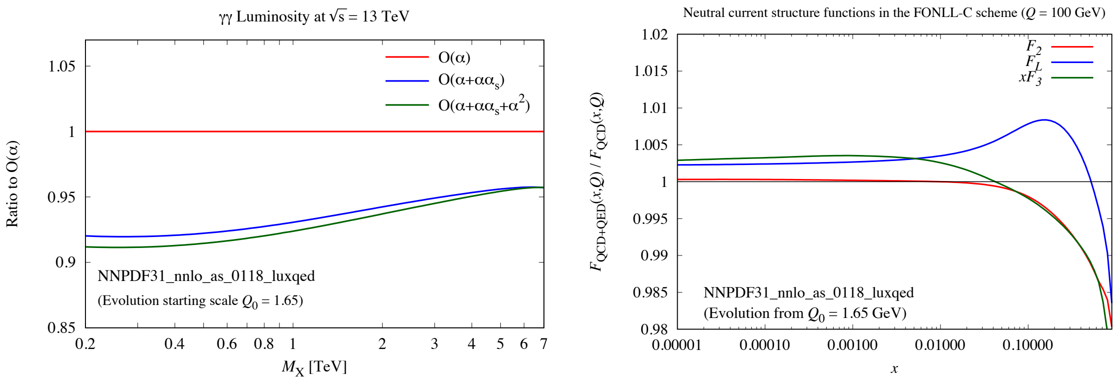

7. *Section 2.4, below (2.2). 'A fraction of its uncertainty' seems a little vague. What fraction is taken?*

   We have replaced "fraction" with the appropriate numerical value used.

8. *Page 8. Ref [116] should be supplemented with 1601.03413, which came before this study (both are referenced in the LUXqed paper).*

   We have added the reference suggested by referee.

9. *Start of section 3.1. Unless I have missed it, the difference between LUXqed 16 and 17 does not seem to be described anywhere in the paper. It would make sense to do this at some point.*

   We have added a sentence in Sect 3.1 describing briefly the differences between the LUXqed16 and LUXqed17 fits:

   "As discussed in Sect. 9.2 of Ref. [42], the LUXqed17 set has a improved evaluation of the photon PDF calculation and of the associated error estimates in comparison to LUXqed16."

10. *Figure 3.1. For the purposes of comparing the cases with only the high Q^2 uncertainties vs. the full case, things are perhaps not presented in the best way. I think it would be helpful to have both cases on the same plot in some way, so that the differences can be seen more directly, but this is not essential.*

    We prefer to leave the comparison as it is now, to avoid cluttering too much the plots shown in the paper.

11. *Page 9, first paragraph. Perhaps it is worth clarifying a little where the dependence on the perturbative order is expected to occur? Surely, at least to first approximation, the only dependence on this comes from the high Q^2 component, as the other LUX components are independent of order?*

    We have extended the discussion of this point and added some additional references.

12. *Fig 3.3 (and 3.4 by implication) and Page 9, second paragraph. The fact that 3.0 photon undershoots the 3.1luxQED photon at low x is surely surprising. In particular, at high scales and low x (i.e. Fig. 3.3. right) the photon is entirely driven by perturbative DGLAP, i.e. in terms of the other partons. Given the compatibility of the 3.0 and 3.1 quark/gluon PDFs, I do not understand how the photon PDFs can look so different in this region. Might it be that the 3.0 photon is calculated using the 2.3 evolution procedure (subsequently corrected)? In any case, given the size of the difference relative to the PDF uncertainties of NNPDF3.0QED the reason for this apparent discrepancy has to be discussed. This all feeds through to Fig. 3.7, where again no explanation for the tension at low mass is discussed. Then again in Section 4, differences in various predictions at lower mass are seen for the same reason, but not discussed.*

    This is an important point that indeed should be clarified, and we thank the referee for raising it.

    First of all, we confirm that in terms of perturbative evolution, the same approach is used in NNPDF3.1luxQED and NNPDF3.0QED, with the only difference of the inclusion in the latter of additional higher order QED corrections. These can already account for some of the differences. To quantify this effect, in the plot below we show the photon PDF $\gamma(x, Q^2)$ evolved from the NNPDF3.1luxQED boundary condition at $Q_0 = 1.65$ GeV using different settings for the DGLAP equations, in particular including only $\mathcal{O}(\alpha)$ terms, as was done in NNPDF3.0QED, and including the full set of QED corrections available, as done for NNPDF3.1luxQED. As can be seen, this can explain up to a factor 5% the differences between NNPDF3.1luxQED and NNPDF3.0QED.

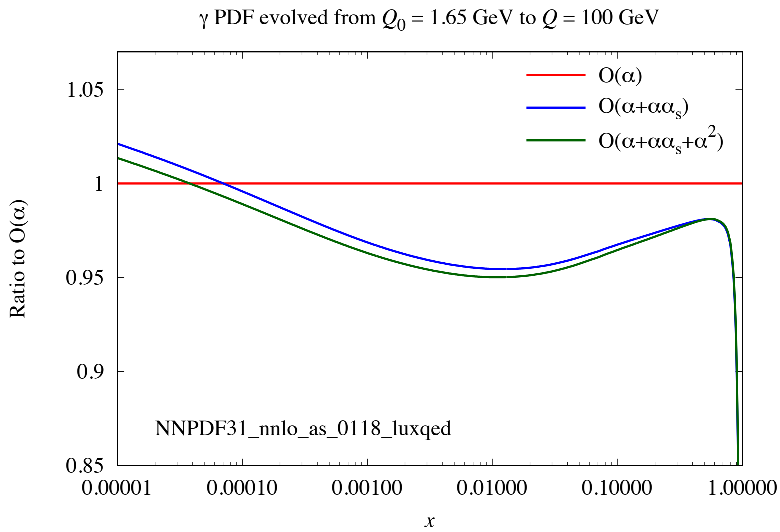

    The other reason for this difference is related to the fitted boundary conditions, namely $\gamma(x, Q_0)$, in NNPDF2.3/3.0QED. As discussed in arXiv:1606.07130, the DGLAP evolution of the NNPDF2.3QED sets was such that the subtraction of the QCD and QED collinear divergences was done separately. This implied the introduction of two different factorization scales, and that the DGLAP evolution with respect to each of them was performed sequentially and independently. This approach, as compared to the more standard procedure, namely the *unified* solution of the QCD+QED DGLAP equations, led to a suppression of the photon PDF at high scales and small values of the Bjorken-x. As one can see in Figure 3 of that reference, taking the same boundary condition and evolving upwards with either the “2.3QED” or “3.0QED” settings leads to differences on the photon PDF $\gamma(x, Q)$ that could be up to a factor of two or three at small-x.

These differences at the level of QED DGLAP evolution in turn affected the resulting boundary condition $\gamma(x, Q_0)$ obtained from the fit. It is however very difficult to quantify by how much $\gamma(x, Q_0)$ would have changed in NNPDF2.3QED at small-x if the unified solution of the QCD+QED DGLAP equations (as in NNPDF3.0QED) would have been used. The reason is that $\gamma(x, Q_0)$ was mostly determined by the Drell-Yan data from low to high invariance masses by means of a non-linear relation.

In this respect, the NNPDF3.0 solution, which was based on taking the NNPDF2.3QED boundary condition $\gamma(x, Q_0)$ and evolving it upwards with APFEL using the unified solution of the QCD+QED DGLAP equations can partially, but certainly not completely, compensate for these differences, specially at small-x.

Therefore, the differences observed between NNPDF3.1luxQED and NNPDF3.0QED can be explained mostly due to the different evolution settings between 2.3QED and 3.0QED, which affected the determination of the boundary condition $\gamma(x, Q_0)$ in NNPDF2.3QED and that were only partly compensated by the NNPDF3.0 fix. The subsequent differences in the evolution settings between 3.0QED and 3.1QED provide a further contribution to this difference.

Following the discussion of Fig. 3.3 we have added a paragraph discussing this point in some detail. This new paragraph reads as follows:

"As shown in Fig. 3.3, for $x \lesssim 10^{-2}$ the NNPDF3.0QED photon undershoots the 3.1luxQED one both at low and at high scales by an amount which is not covered by the PDF uncertainties of the former. There are two main contributions to such difference. First of all, the inclusion of $\mathcal{O}(\alpha^2)$ and $\mathcal{O}(\alpha\alpha_s)$ terms in the DGLAP equations (absent in NNPDF3.0QED), accounts for a difference of up to 5% when the photon PDF is evolved from $Q_0 = 1.65$ GeV to $Q = 100$ GeV (see also Fig. 2.1), explaining part of the discrepancy. The second, and more important reason, is related to the fact that in NNPDF2.3QED the boundary condition $\gamma(x, Q_0)$ was determined from the DIS and Drell-Yan cross-sections using different settings for the QCD+QED evolution equations [39] as compared to those used later to construct NNPDF3.0QED. This partial mismatch then leads to a suppression of the photon PDF at small-x, explaining most of the differences observed in Fig. 3.3."

13. *Fig 3.3. Did the authors intend to take an absolute plot on the left and ratio plot on the right? Given ratios are considered everywhere else it would be more consistent to take that on the left, but clearly this is a minor point.*

    We confirm that this was done on purpose, in order to illustrate the absolute shape of the photon PDFs (rather than the relative one) at low scales.

14. *Section 4. How is the scale of α treated for the coupling of the initial state photons in the matrix element? Historically many people have wrongly used $\alpha(0)$, but as discussed in 1605.04935 and 1705.00598 this is not appropriate for the case of initial state photons with corresponding photon PDFs, even though these are treated as on-shell in the matrix elements; instead $\alpha(\mu_F)$ should be taken. The scale choice should be mentioned, and if the on-shell coupling is used, corrected.*

    The `APPLgrids` generated for the phenomenological studies presented in section 4 have been evaluated at LO in EW theory. In this LO context, α, g_μ, M_Z and M_W are interconnected through the "standard" model in aMG5. Concerning α, we have used the default standard model value in aMG5 2.6.0 which is $\alpha = 1/132.51$. We do not use $\alpha(0)$ anywhere, since α is fixed (no running coupling effects). We have added a sentence in the second paragraph of Section 4 explaining our setup.

We emphasize that our calculation does not include higher-order electroweak corrections, so effects of the running of the QED coupling are beyond our quoted precision. In any case given our choice of $\alpha = 1/132.51$ running coupling effects are most likely a small corrections for the cross-sections discussed in this work. In addition, a more thorough phenomenological study of PI and NLO EW corrections to a range of LHC cross-sections, where effects of the running coupling will be properly taken into account, is ongoing.

15. *Fig. 4.3. Perhaps it is worth emphasising the difference in scale on the y axis relative to the other plots at some point, just for clarity.*

    We have added the clarification requested by the referee.

16. *Page 16, first full paragraph. Typos- should be 300 GeV and M_{ll}.*

    We have fixed these typos.

17. *Page 18, last full paragraph. Again the statement about consistency with respect to 3.0 should be rephrased in light of the discussion above.*

    This has been done, consistently with our answer to point 12) above.