We are grateful to the referee for their comments and for approving our paper for publication, subject to their final recommendations. Indeed, there is some difference in our views (as our exchange has made clear!) but we would like to emphasise that we are in complete agreement with the referee that the LHeC would be an exciting and well–motivated opportunity for investigating proton structure and much else besides.

We reply to the referee's comments in turn below:

In summary, I may also quote from the paper "On the Relation of the LHeC and the LHC" by J. L. Abelleira Fernandez et al., arXiv:1211.5102...

We fully support the sentiment here, and of course agree that the lepton–hadron collider environment is in general theoretically cleaner, providing a strong motivation for the sort of procedure described here. On the other hand, one should be careful not to overstate this: factorization has been robustly proven for Drell–Yan production (which is of course a cornerstone of PDF fits to hadron–hadron collider data) and there are clear phenomenological reasons to believe it holds well in other cases (not least the fact that one can achieve a good description of the extensive global PDF datasets using universal PDFs). So we have very good reasons to believe that global PDF fits provide a robust determination of PDFs, and in our view it is well motivated to asses the impact of the LHeC within this context.

Having said that, we would like to point out again that we do consider the $T = 1$ case, as well as fits to LHeC data alone within precisely the context we believe the referee would prefer. The $T = 1$ comparisons within the global context (albeit with the caveat that this is not completely consistent with the tolerance of the PDF4LHC baseline) are shown in Fig. 4.2. More significantly, the results in section 5 when the baseline priors are reduced by a factor of 100, in fact precisely correspond to a $T = 1$ LHeC–only projection. This point is arguably not made particularly clear in the previous version of the paper, so we have added some discussion of this to the end of the introduction and to section 5 directly. In the latter case we have replaced Figs. 5.4 and 5.5 with results which show the same information as before, but now in a way that indicates the impact on the PDF uncertainties themselves, as well as the ratio of the PDF4LHC to HERAPDF cases (albeit implicitly). Thus our paper shows both the impact of the LHeC within the context of a global fit and within a $T = 1$ LHeC–only fit.

Hence I indeed request that you remove the statement about the 'complementarity of ep and pp data' e.g. remove the last sentence of the abstract and also later in the summary.

We included this sentence as we believed it actually helps build the case for the LHeC (i.e. complementarity is a good thing!), however we are happy to go with the referee's suggestion here and have removed this phrase.

Further on the abstract it is important to realise that the high x potential of the LHeC data vs LHC has not been evaluated at the same footing because the LHeC jet data have not been available. Please add the word 'inclusive' in front of the LHeC data mentioned in the abstract (second last sentence) and elsewhere (e.g. in the summary p13) to make it very clear, in the abstract you correctly state the used data, however, the word "inclusive" will help the reader to understand clearer *the origin when you present the low x and high x data results.*

We see the referee's point here. To be a bit more precise we have added 'inclusive and semi–inclusive' as after all we do include heavy quark and strange production.

Table 2.1 : I see that the author of http://hep.ph.liv.ac.uk/ mklein/lhecdata/datlhecreadme made here the comment of $\eta < 5$, however, if I check the tables of data, the actual eta of usable data is $\eta < 4.4$, anyway, in the context of your cuts it may not matter What I still do not understand, is if you used 1 ab-1 for e+p data or if this is a typo in the table. As Ref. [26] states, e+p luminosity is at most 0.1 ab-1. The availability of positrons in the linac-ring ep configuration is much inferior to electrons. Your results will not depend on this as you find that above a few 10 or so fb-1 the uncertainty is systematics limited.

Thank you for spotting this, and our apologies for missing it. This is a typo- it should say 0.1 ab^{-1} for the positron case, which is what we use in the study. In addition, when checking this we have realised that the quoted heavy quark cases are too low: we in fact assumed 1 ab^{-1} here, as seemed reasonable give the luminosity should be the same as in the inclusive case. We have corrected the table now, though the referee is right that the results will not be too sensitive to this, given one quickly becomes dominated by systematics.

Page 4, paragraph 2.1, line 14: In the current scenario, the charm e+p data would constrain strange like e-p would constrain anti-strange. It would be interesting to see how one may constrain s-sbar from these two measurements. One could probably use the same simulation also for e+ and consider the correlated systematics the same.

This is definitely a good suggestion. In the end we only wanted use those pseudodatasets that were available publicly, but in the future this would be worth pursuing.

To reflect and acknowledge the other view, I recommend that you add the before mentioned Ref. [arXiv:1211.5102] to your paper.

We have added a reference to this in the introduction and conclusion.

There also seems to be a misunderstanding w.r.t. used parameterisation: The parameterisation for LHeC will be determined by the LHeC data and, due to much higher statistics and range, it will surely differ, both from the HERA ansatz, also because ALL PDFs will be determined not just uv, dv, sea, c, b, xg, but as well ubar ne dbar, strange, top and all with much higher precision. It will be a "wonderful world" for QCD and PDFs and I sense there is no disagreement about that expectation with the authors.

Here we would respectfully argue that the misunderstanding is on the side of the referee. We completely agree that the relevant parameterisation required by the actual LHeC data would differ from/need to be more flexible than the HERAPDF case. The issue we discuss in section 5 in this respect is completely specific to the fact that we are performing PDF projections studies, where one must assume a given parameterisation when generating (and fitting/profiling to) the pseudodata. We have tried to clarify this point further by adding a little more discussion to the introduction, which we quote below:

"The slightly subtle issue here is that when one generates pseudodata assuming an underlying PDF parameterisation one is implicitly assuming that it will be possible to describe the actual data with this parameterisation. If this assumption is too strong, that is the underlying parameterisation is too restrictive, then one is liable to overestimate the impact of the data."

Thus exactly as the referee says, one will expect that the very precise LHeC data will demand a more flexible parameterisation than HERAPDF. The point is simply that if one takes this parameterisation (or something closely based on it) when generating LHeC pseudodata then one is implicitly assuming that the HERAPDF parameterisation *will* be sufficient to describe the LHeC data. From a statistical point of view this will lead to an overestimate in the impact on the eventual PDF uncertainties. This is the issue we discuss in detail in section 5.