Response to reviewers' comments on the manuscript "ACE-FTS observations of acetonitrile in the lower stratosphere" by Jeremy J. Harrison and Peter F. Bernath.

We respond to the reviewers' comments, which are in bold text:

Referee #1:

Regarding the MkIV/ACE discrepancy above 25 km. The fact that Kleinboehl et al [2005] fitted pseudo-lines to the PNNL cross-sections is besides the point. Any other extrapolation technique would have had equal if not larger extrapolation uncertainties. So the salient is not the pseudo-lines, it is the limited T/P range of the lab PNNL measurements, which were the only cross-section data available in 2004.
We have incorporated this point into the manuscript.

The chosen window for the ACE retrievals is narrower than that used in Kleinboehl et al. The authors claim that this reduces uncertainties from interfering CH4 and H2O lines, which is true. But there is also a downside to making windows narrower which is that the "continuum level" or "Baseline term" becomes less certain. But this downside is not discussed or quantified. Could the difference in the chosen window widths partly explain the differences between MkIV and ACE at the higher altitudes?
Uncertainty in the baseline term simply contributes to the random error of the retrieval. We have added a sentence to the manuscript to reflect this. However, we do not believe that the difference in window widths contributes any systematic error at high altitudes. As discussed in the paper, errors in CH_3CN spectroscopic data are the most likely explanation for this.

You have a whole paragraph discussing minimum and maximum altitudes for the CH3CN retrieval, but you don't explain why it is necessary to prescribe the altitude range. Why can't you do the retrievals for all altitudes and then look at the errors bars later to decide what the useful altitudes ranges. This way, you would not be prejudging what the H2O or CH3CN profiles should be. At mid/high latitudes you can get days when the tropopause altitude drops to 6km, which would allow CH3CN retrievals to extend much lower, if you let them.
There are a number of reasons why all ACE processing versions have proceeded in this manner. The main reason is that only a small amount of information would be obtained at the expense of the additional processing time involved to extend the retrieval range. A complete reprocessing for the entire mission takes about a year. Processing time is better used elsewhere.

I was surprised to read that a quadratic baseline term was added to the forward model to minimize residuals arising from non-voigt H2O lineshapes. Firstly, I strongly suspect that the main problem in fitting strong H2O lines in the ACE spectra comes from the neglect of the finite FOV. Secondly, if the non-Voigt lineshape is really the issue, why not use the appropriate non-Voigt lineshape, as explained in Boone et al.[2007], rather than fit a non-physical baseline quadratic term?
The non-physical quadratic baseline term improves the retrieval at the lowest altitudes, where the curvature of the baseline due to the saturating H_2O lines is most apparent. In order to use the appropriate non-Voigt lineshape one requires accurate spectroscopic line parameters. There are currently no such parameters available, so we resorted to a quadratic baseline term. Similarly for the field-of-view effect: a quadratic baseline is empirically found to improve the

retrieval. The retrieval will be revisited for the upcoming official ACE processing version 4.0, for which this issue will be looked at again.

In the fig 4 caption, replace "Six MkIV profiles" by "The average of six MkIV profiles"
This has been done.

Referee #2:

GENERAL COMMENTS

Page 3328: Presumably the retrievals don't use an a priori. It would be good to stress this.
All retrievals use an a priori VMR profile of some sort. We assume that the referee is using the term 'a priori' as it refers to the optimal estimation method (OEM). For an ACE v3.0 retrieval, the a priori serves as the initial conditions for the retrieval, which proceeds via a non-linear least-squares fitting technique. We believe that an informed reader will understand this point without any further clarification.

Page 3330, lines 16-18: It would be useful to say a word how the classification to biomass burning plume is made (more than referring to Tereszchuk) and how reliable it is for the purpose of this paper. It would also be useful to show the biomass burning profiles (e.g. as the median or average of these) and see how they differ from the background profiles (what are typical vmr for the biomass burning plumes?). More generally it is unclear to me why the biomass burning observations are not presented in the paper. They would give the paper more weight.
We have provided more information on the classification of biomass burning plumes, as requested.
When discussing satellite observations of biomass burning plumes, it is more appropriate to consider individual profiles that can be associated with a particular plume. As discussed in the manuscript, the retrieved ACE CH_3CN VMRs for a single profile have large random errors, comparable to the magnitudes of the VMRs. In order to reduce the random error contribution to the profile, we need to take averages of many profiles. Due to the limited global coverage of the ACE-FTS instrument (a maximum of 30 measurements per day, and most of these at high latitudes where biomass burning is not so common), we would need to average profiles associated with plumes of different types and ages, and over different time periods. One must also note that the retrieval is essentially confined to the lower stratosphere; we would expect more variability due to biomass burning in the upper troposphere. For these reasons, we chose to focus specifically on the background CH_3CN concentrations, for which we believe the data are best suited.

Page 3331 and Figure 2. The profiles are built from a set of profiles in latitudes bins and for several years and this is justified by the fact that the inter-annual variability is small. But how much is it? And how does it compare to the latitudinal (or hemispheric) variability? In fact all profiles shown in Figure 2 seem to be similar within the retrieval error. The question of what is significant and what is not should be addressed more critically. Relating to the above, one would also like to know if the biomass burning activity does not cause larger variability (in space and time). Why is it not discussed here?

Given the large random errors in the retrieved VMRs, it is not a simple task to determine the inter-annual variability, particularly because the ACE-FTS only takes a maximum of 30 measurements per day. Annual latitude-band plots indicate that this variability is well within the error bars shown in Figure 2; this is the best we can conclude from the data. For this reason it was decided to plot all data together, thus improving the profile statistics.
The profile plots in Figure 2 do reveal a number of significant features. We have improved the explanation surrounding Figure 2 to clarify this.
Biomass burning does of course create variability, more so in the upper troposphere, however, as discussed above, the poor global coverage of the ACE-FTS and the large random error contribution to the VMR profiles make it difficult to quantify this variability. For the reasons discussed above, we only focus on background concentrations of acetonitrile.

Page 3331, Figure 3: The text around Figure 3 is very descriptive. Is there an explanation for the seasonal variability? In the Northern hemisphere the enhancement are above 50 in winter. Is this explained by chemistry? And if there is a chemistry effect, would it affect the lifetime estimated in section 7 (winter vs summer). The tropopause altitude should be indicated in Figure 3 (or are the measurements all for sure in the stratosphere?)
We believe the seasonal variability is related to biomass burning contributions coupled with the six month atmospheric lifetime. We have expanded the text around Figure 3. Note that although it appears that the NH enhancement occurs at latitudes above 50°N, this is probably an optical illusion since the retrieval extends to lower altitudes above 50°N. The only chemistry involved is the biomass combustion. We do not believe that the seasonal enhancements affect the estimated lifetime within the uncertainty limit specified. The enhancements are largest only for the lowest altitude levels (in Figure 3); these will only create some additional scatter at one end of the linear plot in Figure 5.
Most of the measurements in Figure 3 are above the tropopause; it is expected that measurements at the lowest few kilometres in the tropics are below the tropopause. It must be realised that the data are averaged and that the tropopause can vary from day to day. We have incorporated these points into the text instead of modifying the figure.

Page 3333, last paragraph before section 7. The MLS v3.3 CH3CN data seem to be publically available. I understand that they should be used with caution but considering that they are the only correlative dataset (matching the time period covered here with ACE-FTS) the authors should provide better comparisons (a least one Figure) and supporting discussions.
We agree that in normal circumstances there should be a detailed comparison with MLS data. Preliminary work was performed in this direction, however the agreement was poor; it is clear to us that there is something wrong with the MLS dataset. It is not the purpose of this paper to criticise MLS, so it was decided not to proceed with a detailed comparison. We hope that our work will provide motivation for the MLS science team to improve their retrieval. Below is a figure comparing the median ACE profile for |latitude| > 60° (where the ACE error bars tend to be smallest) with the mean MLS profile for measurements taken in the same latitude range on 17 June 2011.

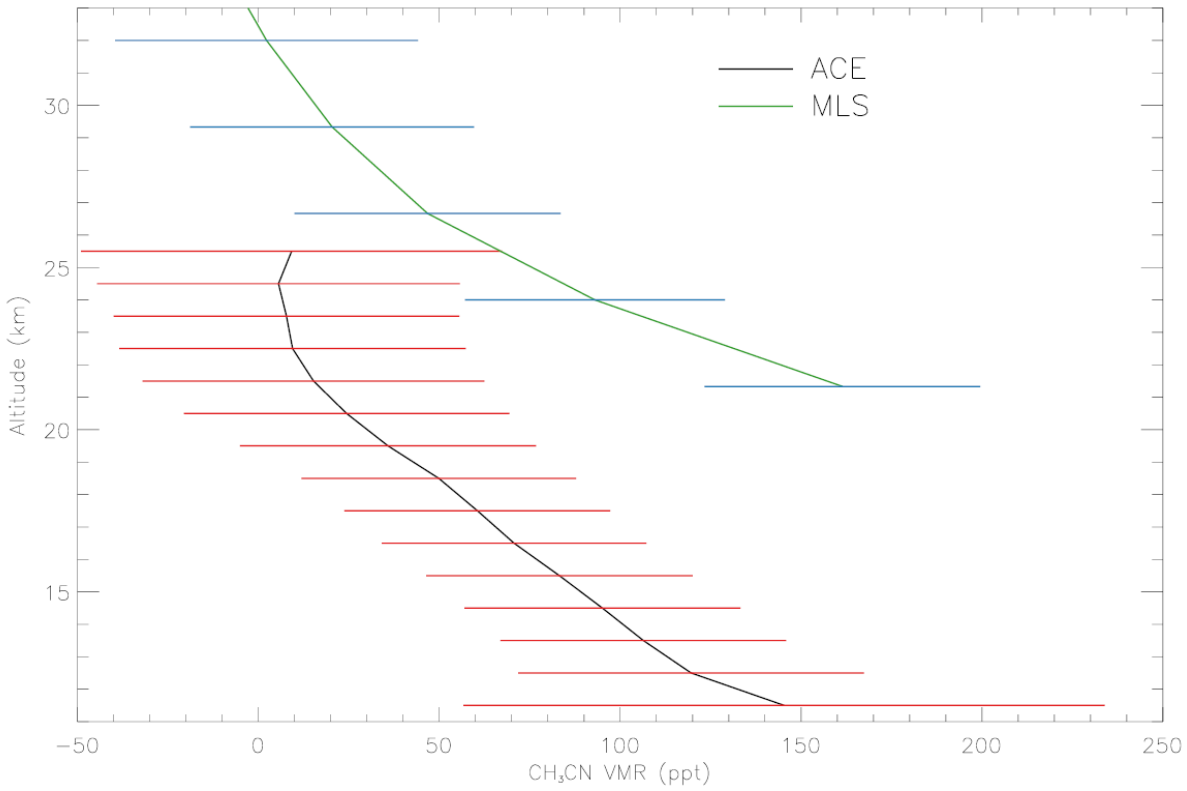

Page 3335, before section 8. The lifetime is about 5 times larger than previous estimate. The statement "this work supports the view that acetonitrile is long-lived" falls short. A tentative explanation and/or a short discussion would be needed.
It is not completely clear to us what the objection is to this statement, however it has now been removed.

MINOR AND TECHNICAL CORRECTIONS
Page 3324, line 18: Distinction should be made between the tropospheric lifetime and the stratospheric lifetime discussed later in the paper. The 6 months is presumably a tropospheric lifetime.
The figure of 6 months is actually the total atmospheric lifetime, although since the stratospheric lifetime is over 100 times larger, the total lifetime is approximately the same as the tropospheric lifetime. This distinction has been clarified in the manuscript.

Page 3325. Equation 1 is not discussed elsewhere; R1 could be removed
R1 has been removed as suggested.

Page 3326, line 24. Remove "between 1993 and 2004", which is already mentioned above.
This has been done.

Page 3331, line 4: "The 18 latitudes bins from -90 to +90"
This has been done.

Page 3331, line 12: It is unclear what this Gaussian distribution means. Is it useful at all?
One would expect VMRs at each altitude in each latitude bin to be normally distributed. This sentence just confirms this fact.

Page 3325, lines 12 to 21. This paragraph is a bit confusing. It relates to the tropospheric sources and sinks and gives details that are not relevant for the paper. I would suggest shortening.
We have simplified this paragraph as requested.

Reference: Tereszchuk 2013 should be 2012 (as also quoted in the text Page 3329, line 27)
This has now been published in ACP, so the reference has been updated.