Response to Reviewer 2

We would like to thank the referee for the helpful review and the constructive comments. His/her comments are given below in black and our responses in blue.

General comments:

As the developer of an LTE radiation scheme used in a GCM, my review is from the point of view of a potential user of the code who would want to improve representation of the upper atmosphere in their code. There is certainly a strong case for a revised non-LTE parameterisation that can handle much larger concentrations of CO2, and therefore the contribution from this paper is welcome. But as I am not an expert on non-LTE effects I can't comment on the scientific details of this particular parameterisation. Naturally I would expect the code to be available at the time of submission to GMD and I expect the authors to rectify this.
Reply: We appreciate your valuable suggestions even if not being an expert on non-LTE.
The referee is fully correct about the availability of the code. We were not aware of the need for its availability during the review process and thought of providing it during the review process. We will provide it in its final version as a Fortran 90 code with the revised version of the manuscript. In the meantime, the parameterization is now available provisionally as a Python routine, please see the reply to the Editor comment 1 (https://zenodo.org/doi/10.5281/zenodo.10547026).

The algorithm is presented as a complete radiation code, but the LTE part described in section 5.1 is quite crude for tropospheric radiative transfer since it doesn't represent scattering or cloud overlap and heterogeneity. I imagine that most GCM developers (like myself) would want to use their existing radiation scheme in place of the algorithm described in section 5.1 for the LTE region from 0 to 70 km, and then to add on NLTE regions 1-4 described in section 5.2-5.4. It would therefore be helpful to describe how to do this:
Reply: We did not make it clear in the text. It is not intended to use this radiation scheme in the LTE region, at least not below 50 km. It was not our purpose to develop the parameterization for that region. Note that it includes only CO2 and hence it is not designed to be used in the LTE region where other species contribute to the cooling/heating.
We will clarify this in the text and the documentation of the code.
The suggestion is that users use their own LTE radiation scheme at low altitudes, in the LTE region and then switch to the non-LTE region with the cooling rates provided by this parameterization. A safe switching region could be between 50 and 60 km. Still, this parameterization can be used with a reasonable accuracy down to the lower stratosphere, although it should be kept in mind that accounts only for CO2 cooling/heating.

Can I simply replace the cooling rates from my code with the cooling rates from the non-LTE parameterisation above 70 km.
Reply: Exactly, or from even above 50 km. A switching region is advisable to be used.

Do I need to provide any upwelling fluxes at 70 km to feed the treatment of the regions above?
Reply: No, the parameterization calculates internally the entire cooling rate profile (with its 'own' upwelling flux from the LTE region). However, the input file (with pressure, CO2, T, etc.) should cover the entire atmospheric range, e.g., from the surface up to at least x=log(1e3/p(hPa))=13.5, i.e. p=1e-3 hPa (or ~92 km) in order to properly calculate internally the upwelling flux. If some layers at the bottom of the atmosphere are omitted, the parameterization still works but the cooling rates are not fully correct. Hence that is not recommended.
A surface temperature is also an option to be imputed. If not given, it will be taken as the value of the lower altitude of the input file.

Again, all of this is already documented in the Readme file of the version in the repository.

The description of how one region is coupled to the one above is not clearly described in the text.
Reply: That's correct. We are giving recommendations above the altitude range for the use of the parameterization, as mentioned above. But the user only needs to consider two regions: the lower one, in LTE, and then this parameterization on the region above.

Is the vertical resolution of the parameterisation fixed or can I run it on my own vertical grid?
Reply: The user can provide its own vertical grid in a specified input file which should contain the pressure (hPa), temperature (K), VMRs of CO2, O, O2, and N2. Then, the parameterization interpolates those parameters onto its internal grid (described in the manuscript). This is detailed now in the readme file of the code.

Do I need to simulate all regions or can I omit the uppermost 1, 2 or 3 regions if I am not interested in modelling temperatures above a particular height?
Reply: Again this is an interesting question. In principle, the users should input the parameterization from the surface up to the upper boundary of interest. This upper boundary cannot be lower than ~92 km (1e-3 mb). However, it is recommended to extend the upper limit up to ~120 km since exchange or radiation takes place between the upper mesosphere and the lower thermosphere.

Has the parameterisation been coded in Fortran or C (needed for a GCM) or in an interpreted language like Python? What is its computational cost? How large is the dataset that needs to be read in (e.g. the various look-up tables that are used to populate the Curtis matrices)?
Reply: The current (provisional) version in the repository is in Python but we will make it available in Fortran with the revised version of the manuscript.
That version of Python is very slow. We will provide the computational cost when translated to Fortran.
There will not be an external dataset for reading. The Fortran routine will contain all the needed coefficients as parameter variables.

SPECIFIC COMMENTS
1. Abstract: it would help to mention the magnitude of LTE heating rate errors (i.e. LTE minus non-LTE), in order to put into perspective the magnitude of the errors in the parameterisation of non-LTE effects (i.e. parameterised non-LTE minus accurate non-LTE).
Reply: We agree. They very much depend on the altitude but in general they are way smaller. We will include a general statement about these differences.

2. There is an excessive number of figures, and they are sometimes referenced out of order. My suggestion is to select one (or at most two) of the test atmospheres and then when you have a 6-panel figure with a separate panel for each atmosphere, reduce it to just one panel in the main body of the text, and then combine figures together when it is useful to have panels next to each other. For example, combine Figs. 6a and 7 into a 2-panel figure, since they show the same thing for the same atmosphere, just on a different scale. Then put the other five atmospheres in Supplementary Material (not in appendices) and refer to them sparingly. I suggest that Fig. 1 is removed as it is not needed, and Fig. 2 is replaced by Fig. 11 (the latter shows the same as the former but with more information). There are plenty of other opportunities for the authors to cut down the number of figures.
Reply: We agree. The other referee has made a similar comment. The reason for showing the results for the six p-T profiles is that frequently, the effects of the discussed parameter depend very much on the temperature profile. Nevertheless, we take the suggestion. We will make a drastic reduction in the number of figures/panels in the main text. They will be reduced from 28

to 19 (and some of them with a reduced number of panels); and the number of Figs. in the appendix will be reduced from 20 to 8. 21 figures will be moved to a Supplement.

Please see the response to the first general point of the other referee for all the detailed actions taken regarding this point. Further, Figs. 2 and 11 are not merged because, if we do, we will not see the differences in the temperatures in the mesosphere and lower thermosphere, precisely the region where the major difficulty of the parameterization resides. About "and refer to them sparingly", we will try to make the reading smooth but in general we prefer to point the reader to the figure in question so he/she can better understand the argument. He/she always cannot look at the figures if unnecessary.

3. Lines 102, 200 and elsewhere - please refer to profiles #3 and #6 as "present day" and "4x pre-industrial" so the reader doesn't need to flip back to remind themselves what these numbers refer to, and similarly for the other hash-numbered profiles. Also, Fig. 9 could state the factor multiplying the pre-industrial figure surface value for each point somewhere on the figure.
Reply: Definitely, we will do that. It will make the reading much easier. About Fig. 9, it will be moved to the Appendix but we will include in the caption the factor multiplying the pre-industrial figure surface value for each point and will make a reference to Fig. 2 (the CO2 vmr profiles).

4. Line 219: "pointing" -> "role in determining"?
Reply: Thank you. That was a leftover of a previous version. It will be corrected.

5. The equations are largely taken from Fomichev et al. (1998), but it would be useful to improve clarity in several places. For example, Equation 1 is simply a matrix-vector multiplication - why not present it as such? The equations for "a" and "b" on lines 264 and 265 are presented as if they define the element the Curtis matrix (via the equation on line 261), but they themselves contains the elements of the Curtis matrix on the right hand side as a scaling for the terms involving the band strengths. So one is left wondering how the actual value of the elements of the Curtis matrix are determined.
Reply: "Equation 1 is simply a matrix-vector multiplication - why not present it as such? " Yes, writing it in a matrix formulation will be more concise (and clean), but we prefer to write it explicitly, particularly because the summation goes over different altitude ranges and it is easily written in this form. If using the matrix notation we would need to define different matrices (symbols) for the different altitude ranges.

The Curtis matrix is actually embedded in $\mathcal{A}^t_{i,j}(\nu)$, see lines 253-254.
a and **b** are scaled Curtis matrices.

6. Line 404: remind the reader that by "both" parameterisations you mean Fomichev's original one, and the one in this paper.
Reply: Sure. We will do that.