We thank the reviewer for their additional comments on this manuscript. Please find our responses given in the blue font below (note all references to line numbers refer to the track changes version of the revised manuscript).

Lines 98-99: Given that the O3 perturbation is bounded by the approximated tropopause [i.e., the linearly varying tropopause (from 100 hPa at the equator to 300 hPa at the poles)], does this imply that there are O3 perturbations within the actual troposphere? Alternatively, does this tropopause setting lead to O3 changes within the upper troposphere and subsequently affect high clouds? It would be helpful to provide the global-mean or zonal-mean vertical cloud fraction climatology distributions for the different control experiments (standard, 0.5x, and 1.5x Ozone) and the corresponding cloud fraction changes to 4xCO2. Additionally, it would be beneficial to include both the all-sky and clear-sky 4xCO2 IRF to evaluate whether the cloud masking value changes with different O3 settings.

We have now included zonal-mean cloud fraction distributions for each case (base-state and 4xCO_2) in supplementary Figure S3 and these are referred to in the main text (see lines 313-316). From these plots we see that the ozone experiments do lead to changes in high cloud fraction (decreasing in 'Strat O_3x1.5' and increasing in 'Strat O_3x0.5'). Though this appears to occur largely below the approximated tropopause level and the increase/decrease in ozone (see Figure S3a). We have also included all-sky and clear-sky (SW and LW) IRF values in Table S3 of the supplementary and we (briefly) refer to the cloud-masking effect of each experiment on lines 316-320. We hope that the inclusion of more cloud details now helps the reader to understand the role that cloud changes play on radiative fluxes in each experiment.

Lines 107-108: Can the authors clarify what is meant by "12-month climatology"? Is this a multi-year mean climatology (e.g., 25 or 30 years), or does it specifically refer to "the first 12 months of its output," as mentioned at the end of the sentence?

This referred specifically to the first 12 months of output as mentioned at the end of the sentence, but the text has now been re-worded to avoid any confusion. Hopefully this is now clearer. It now reads as:

"It is used here to perform two sets of radiative transfer calculations for each experiment listed in Table S1; a baseline (control) simulation and a perturbed (4xCO_2) simulation, which are both run using a year's worth of climatology from the corresponding ERF control integration (i.e., the first 12 months of its output)."

Lines 127-129: Given the definition of ERF as the net downward radiative flux difference, the cloud radiative effect change (dCRE) should correspond to the difference between all-sky and clear-sky ERF. As such, there should be no negative sign preceding d_cre_lw and d_cre_sw in the adjusted cloud radiative response section (as noted in response to Reviewer #1).

Thanks for highlighting this, we agree there should be no negative sign before d_cre_lw and d_cre_sw, apologies for not spotting this in the response to Reviewer 1. Looking back at our python script, I see that we initially calculate LW CRE (d_cre_lw) and SW CRE (d_cre_sw) as clear-sky minus all-sky, and the minus sign was added later in the script to correct this as all-sky minus clear-sky when coding the equation for dLW_cloud and dSW_cloud. These equations were then copied in the response to Reviewer 1.

Lines 129-132, 164-169, and 241-242: Since the radiative kernel method is now being used for cloud adjustment calculations, I recommend rephrasing these sentences for clarity. Additionally, because the revised method has resulted in non-closure in the decomposition, it would be helpful to provide the residual term values for all decomposition plots in the manuscript. While the original radiative kernel method, the method used by Smith et al. (2020) (including APRP and PRP for cloud adjustment), and the previous version of this manuscript had no residual terms, residuals are apparent in the current version. Furthermore, could the authors provide comparisons of dCRE, cloud masking of IRF, and cloud masking of other adjustments for these 4xCO2 simulations? It would be good to understand why the cloud adjustment values are different between previous and current version of decompositions.

We have now updated the text, thanks for highlighting this oversight (see lines 130-135). We have also included the residual term in Figure 1 and Figure 3 and made reference to this in the text (see lines 322-325).

We have also provided comparison of dCRE and the cloud masking of IRF and adjustments as figures in the supplementary and we refer to these plots in the main text (see lines 309-320). Now that we calculate the cloud adjustment with the use of kernels, residual term ϵ is not aliased into our values and we can compare the relative contribution of dCRE and cloud masking in each experiment.

Lines 132-138: Do the ERF values reported in the manuscript represent land-warming-corrected ERF?

No, the ERF values aren't corrected for land surface warming and so don't represent land-warming corrected ERF. We have now added the following text on lines 141-143 to clarify this:

"Subtracting A_{T_S} from the ERF could provide a land surface warming corrected forcing (following Smith et al., 2020a), however we do not calculate this here. Instead, we report the magnitude of A_{T_S} to inspect any change in its value with each O_3 experiment."

Figure Caption of Figure 1: Note that Smith et al. (2020a) calculated non-cloud adjustments using the radiative kernel method and cloud adjustments using APRP and PRP for SW and LW, respectively.

Thanks for pointing out that this needed to be clarified in the caption. The caption text has now been updated.

Lines 181-182: The statement "whereby geopotential height is used as an approximation of geometric height on model pressure levels in the control integration" seems unnecessary and could be omitted for conciseness.

We agree, this has now been deleted from the text.

Lines 299-300 and 309-311: It would be helpful to explain why stratospheric O3 changes do not significantly affect the magnitude of ERF, despite the significant differences shown in the IRF. Is there a specific adjustment or residual term offsetting the IRF's contribution? If not, why does the ERF fail to reflect the difference observed in the IRF and SARF?

We agree that this would be helpful and we have added text to discuss these points (lines 302-325). Here, we specifically highlight the strengthening/offsetting impact of the adjustments and the residual term to explain to the reader why the ERF fails to reflect the differences seen in the IRF. We also further discuss the stratospheric temperature adjustment (see lines 260-270) to give the reader more background on the nature of this important adjustment.

Lines 311-314: Given the secondary contribution from the spectral overlap of CO2 and O3 to the magnitude of IRF, as demonstrated in the revision, could the authors highlight the dominant role of stratospheric temperature dependence? The current phrasing might imply that the two mechanisms contribute comparably.

We have now clarified this in the conclusion (see lines 347-349).

"Instead, these experiments demonstrate a dominant impact on the magnitude of IRF, primarily due to the impact on base-state stratospheric temperature with an ancillary effect from spectral overlap of CO_2 and O_3".

And also in the abstract (see lines 16-18).

"These experiments impact the IRF primarily due to the influence of base-state stratospheric temperature on the emission of outgoing longwave radiation, with the spectral overlap of CO2 and O3 playing a subsidiary role."

General Comment: Could the authors clarify why the albedo and water vapor adjustment values differ by eye between the previous and current versions in the tracked document?

Apologies for not having highlighted this in our previous revision. The albedo adjustment values differ because we originally opted to use the 'asdir' and 'asdif' (i.e. direct and diffuse surface albedo) NorESM2 output fields in our calculation (which are used in the construction of the albedo kernel, see: https://essd.copernicus.org/articles/10/317/2018/). However, we realised that the user is meant to use the shortwave surface radiation fields to calculate surface albedo (FSDS and FSNS), so we updated our calculations accordingly. The H_2O adjustment differs for a similar reason, whereby we updated our method to be fully consistent with Pendergrass et al. (2018, https://github.com/apendergrass/cam5-kernels/tree/master). A lot has been learnt throughout this paper regarding the use of kernels, apologies for unintentionally excluding this explanation before.