GMD: Response to reviewers

July 13, 2026

We thank the reviewers for their careful and detailed feedback on our manuscript. Below are our responses to the reviewers' questions and concerns. We have incorporated their suggestions into the manuscript where appropriate and believe that the updated manuscript is much improved over the original submission.

Additionally, we would like to note that the original submission contained a bug in the E3SMv2 source code. Specifically, the regularization of $\hat{C}_{\mathrm{HN}}$ described in Section 3.3 was accidentally applied to the CTRL simulation (see L491-497 of `/share/util/shr_flux_mod.F90` in the original release). This bug has been corrected and a new release of the source code has been issued, with the Zenodo citation in the revised manuscript being updated accordingly. Figures 14, 15, C2, and E1 have been corrected to reflect the new results of the CTRL simulation as well.

1 Reviewer 1

1. "When instantaneous daily values of inputs for the flux scheme are saved, is that done at the same simulated time of day? Is it 0Z? And for all 10 years? You might miss some meteorological conditions since half the planet is always night and half always day if always sampling at 0Z but I don't think that would change the results."

   Response: Yes, instantaneous outputs are written once per day at the same simulation time (0Z). We have done an additional control simulation with outputs written four times per day (0Z, 6Z, 12Z, and 18Z) and regenerated figures 3 and 12, which remain largely unchanged.

2. "You should mention how many sets of values were saved. (10 * 365 * # of ocean grid points) total."

   Response: The number of ocean grid points varies at each time step (due to melting and formation of sea ice). In the control simulation (CTRL), the minimum number of ocean grid points is 12 477 and the maximum is 12 979. The total number of instantaneous values in CTRL over the entire ten years is 183 589 464 assuming four-times-daily archiving. We have added this information to the revised manuscript in L240-244.

3. "Was every set subject to the offline calculation mentioned in 3.2?"

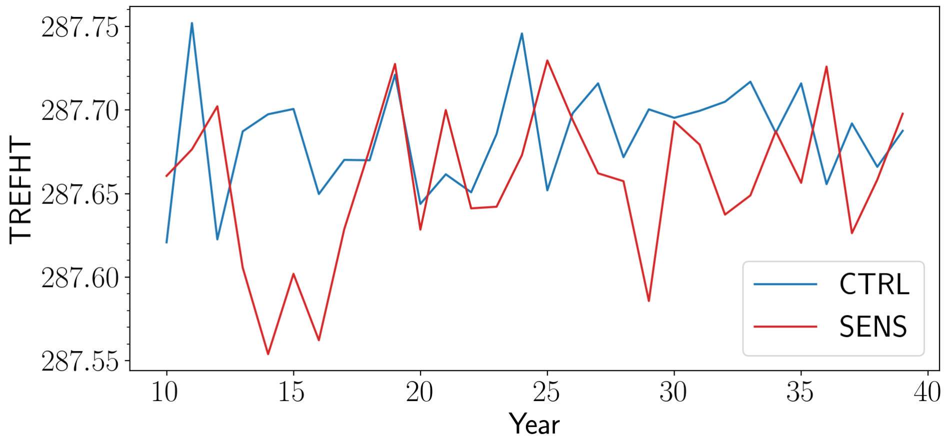

Figure 1: Global annual average TREFHT from 2010-2039 from CTRL and SENS simulations.

Response: The offline calculation described in Section 3.2 was applied to the control dataset only (CTRL) in order to ascertain where and how often the iteration fails to converge, and also test the proposed regularization. It was not applied to the test dataset (SENS) since this dataset already has the proposed regularization applied directly in E3SMv2.

4. "In section 5, results are presented using a 10-year mean calculated from a 10-year simulation. Even in an F-case, there will be some adjustment time because of the land surface. You should look at a time series of global annual average surface air temperature to see how quickly it reaches a quasi-steady state. You may have to discard the first year and use 9-year averages. If a large adjustment to the new scheme is in the first year of SENS, it would distort the tests of significance."

   Response: We have added additional simulation years to the original tests and plotted (1) global annual average surface air temperature (TREFHT) from 2010-2039 (see Figure 1 below) as well as redone Figure 14 from the manuscript (see Figure 2 here). From the surface air temperature in Figure 1, we do not see much evidence of spin-up in the first ten years.

5. "In equation A11, should it be $C_{\mathrm{H}}(U/u_*)$ and $C_{\mathrm{E}}(U/u_*)$? Something isn't right if $\sqrt{C_{\mathrm{D}}} = u_*/U$."

   Response: You are correct. Thank you for catching this error. We have revised Eq. (A11) to rectify this error.

2 Reviewer 2

1. On the specificity and framing

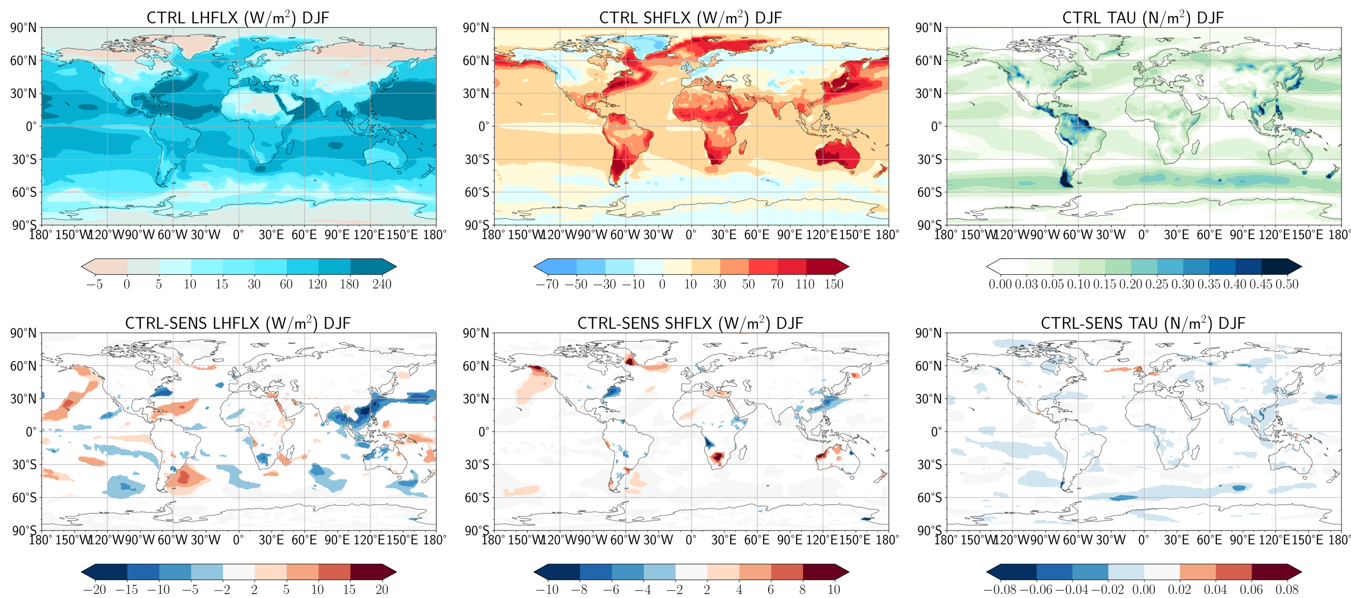

Figure 2: (Top row) CTRL averages of LHFLX, SHFLX, and TAU in DJF. (Bottom row) Absolute difference in LHFLX, SHFLX, and TAU between CTRL and SENS simulations from 2010-2039.

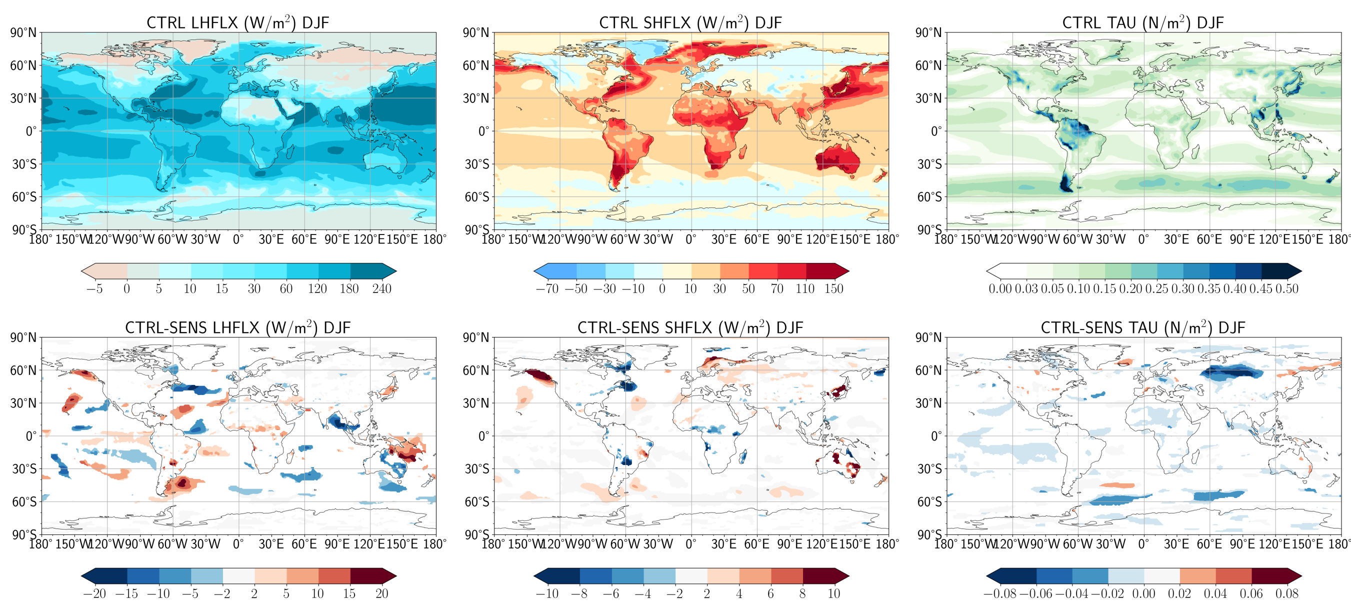

Figure 3: (Top row) CTRL averages of LHFLX, SHFLX, and TAU in DJF. (Bottom row) Absolute difference in LHFLX, SHFLX, and TAU between CTRL and SENS simulations from 2010-2019.

- "In my view, the manuscript should be framed as an in-depth mathematical analysis of one specific parametrisation in one specific implementation. At present, it is presented more broadly than that, including in the title. Alternative air–sea flux parametrisations exist, most notably COARE, but they are treated as a minor detail. A clear example appears at the start of Section 2: Large and Pond (1982) is introduced, and the authors note that the conclusions may also apply to non-E3SM models that use it, but they do not clearly acknowledge most models use different, more modern parametrisations. My suspicion is that many of the conclusions would extend more broadly, but the manuscript should state explicitly that the present investigation was carried out for only one parametrisation."

  Response: We thank the reviewer for this insightful comment. We agree that the title of the original manuscript was more general than the scope of the subsequent analysis and results. We have revised the title to be more specific. We feel that the introduction in the original manuscript starting at L44 already notes that the focus of this work is on a single parameterization, but we have strengthened the language here to note that we are only considering the specific implementation in E3SMv2. We have also added an additional paragraph (L59) noting parameterizations used in other ESMs and that E3SMv2 actually has additional parameterizations implemented, although Large and Pond is still the default (in E3SMv3 and v4 as well).

2. Overlooked literature: "I believe several key references have been overlooked, namely

   - Louis (1979) presents a direct, non-iterative method for computing air-sea fluxes."
     Response: We have added this reference along with Kara et al. (2000, 2005) at L34 of the revised manuscript.
   - "Large (2006) arguably provides a more up-to-date reference for the CORE bulk formulas, following up Large and Pond, 1982."

     Response: We have mentioned that Large (2006) update Large and Pond (1982), but Large and Pond (1982) is still the parameterization used by E3SM as well as CESM and several other models derived from CESM.
   - "Brodeau et al. (2017) examine the implications of several assumptions used in bulk formulas. Had the manuscript used the generic AeroBulk software, which emulates several such formulas, its conclusions might have been stronger."

     Response: We have mentioned Brodeau et al. (2017) at L38 of the revised manuscript where we felt examining assumptions of the bulk formulae fit with the existing discussion of sensitivity to choice of parameterization. More generally, our manuscript seeks to (1) introduce the concept of well-posedness of turbulent flux parameterizations to the general Earth system modeling community and (2) demonstrate how these issues might be addressed and what the result impact on model results is, using the Large and Pond parameterization in E3SM as a proof of concept. Our hope is that the community will extend our work here to other parameterizations and algorithms, perhaps using the AeroBulk software.

The original manuscript mentions possible issues with the other, non-default parameterizations in E3SM (University of Arizona and COARE) with the intent to inspire such future work.

- "Torres et al. (2019) briefly discuss the impact of thresholds and limiters on the stability parameter ζ (see the first paragraph of their Section 4). Although this is not the focus of that paper, such details are rarely treated explicitly in the literature, as the present manuscript itself notes.

  Response: We mention Torres et al. (2019) at L82 of the revised manuscript.

- Kara et al., (2000), (2005); Pelletier et al. (2018), in different contexts, both propose regularised versions of bulk air-sea formulations. Our 2018 paper in particular shares several concerns with the present submission, even if the proposed remedies differ; see, for example, Section 2.4."

  Response: We have added a new paragraph (L94) in the revised manuscript which covers some of the missing references regarding regularization of discontinuities in turbulent flux parameterizations. In particular, we note the polynomial regularizations in Pelletier, et al (2018) and Kara et al (2000, 2005), as well as the discussion of limiters destroying C^1 continuity and lack of accuracy with low number of fixed point iterations in Pelletier, et al (2018).

- "Schüller et al., (2025) address a different topic (coupling and convergence) but the broader perspective is strikingly similar. The authors analyse convergence in a simplified air-sea coupling problem and regularise selected parameterisations (convection and sea-ice albedo in their case) to improve regularity and convergence. This would be worth citing, for example in the conclusion."

  Response: While Schüller et al. (2025) indeed discusses a different topic, their conclusion that discontinuous sea ice albedo parameterizations can lead to lack of convergence of Schwarz waveform relaxation (SWR) methods is relevant to our current work. We have added this reference to the introduction in the paragraph beginning at L94.

3. "Continuity as a necessary cause for fix-point existence"

   - "Around lines 215–222, the logic is somewhat convoluted. In my view, the authors' point only holds under additional assumptions about the asymptotic behaviour of $x \mapsto f(x, \boldsymbol{\xi})$. The example based on Eq. 18 is therefore somewhat misleading. The manuscript partly acknowledges this by noting that some discontinuous mappings do have fixed points, but some continuous mappings also have none, for example $x \mapsto x + 1$. There are many fixed-point theorems, and I suspect that applying one rigorously to these bulk closures would be beyond the scope of the paper. Intuitively, continuity and fixed-point existence do seem related in the case of bulk closures, but the manuscript does not yet provide convincing mathematical evidence for that claim. This a difficult issue. I therefore think the tone should be more cautious here. Continuity is likely a desirable property, but until its necessity is established, that remains an informed expectation rather than a demonstrated result."

Response: We thank the reviewer for pointing out that this section is not precise enough. We have added an introductory line discussing under what scenarios one might have guarantees of convergence to a unique fixed point, and we note that most theorems like the Banach Fixed Point Theorem only provide sufficient but not necessary conditions for this convergence. We note that due to the singularity in ζ, its derivatives are unbounded and hence, derivatives of $\boldsymbol{f}$ are unbounded and thus, $\boldsymbol{f}$ will typically not satisfy strong assumptions like the contraction mapping property required for the aforementioned theorem. We have softened the language in these lines and noted explicitly that we are unable to rigorously prove lack of solution existence and, absent such a proof, rely on the graphical argument in the original manuscript alongside numerical experiments to demonstrate that lack of continuity appears to be an issue under certain meteorological conditions for this particular parameterization.

We also feel that the additional mentions of fixed point theorems help address the reviewer's comment on the relationship between continuity and fixed point existence. Specifically, the stronger Lipschitz continuity with constant < 1 (i.e. contraction mapping property) is sufficient for existence (and uniqueness). Regular continuity is not sufficient, as the reviewer pointed out in their example of $x \mapsto x + 1$.

4. "On the efficiency of the "adaptive" stability limiter"

- "Scientifically, adding a limiter means modifying the problem being solved. I would encourage the authors to state this more plainly in places where the manuscript currently alludes to it indirectly"

  Response: We have made this more explicit when first introducing the limiter (L219 of the revised manuscript), and then remind the readers of this in Section 4 (L446 of the revised manuscript).

- "The discussion on the $\zeta_{\max}$ adaptive limiter is difficult to follow. I would suggest moving part of the technical detail out of the main text and into an appendix. Sections 4.1.1 and 4.1.2 could likely be reduced to a short summary in the main body: depending on the precise parameter choices, and on how the $\zeta_{\max} \to 0$ limit is approached, existence and/or uniqueness may be lost. The detailed examples would be better placed in an appendix."

  Response: We appreciate the reviewer's insights here and see that the denser mathematical arguments in Section 4.1.1 can be confusing to follow. We have combined Sections 4.1.1 and 4.1.2 and moved much of the content of Section 4.1.1 to a new appendix.

- "More importantly, I remain somewhat sceptical about the $\zeta_{\max}$ adaptive procedure. By the authors' own account, the method still falls back, as a last resort, on the hard limiter the $\zeta_{\max} = 10$. The statement in lines 492–494 is therefore not adequately supported. As the manuscript rightly notes in the introduction, many parametrisations are developed without explicit attention to mathematical well-posedness. The need to revert to the default hard limiter in some cases shows that the problem remains ill-posed, but it does not establish why. It may be that the underlying physical assumptions fail in those regimes, but it is equally possible that the physics remains acceptable while the algorithmic formulation breaks down. The manuscript seems to favour the former interpretation without providing evidence, and I suspect such evidence would be difficult to provide. In my view, the fallback is an acceptable limitation and is consistent with the broader message of the paper: these parametrisations are not always well posed and often require ad hoc fixes to behave robustly."

Response: We thank the reviewer for their helpful comments on this matter. It is certainly true that the fallback to the hard limiter is not desirable. We expanded on this fallback scenario and its occurrences in the revised manuscript. In particular, the fallback is intended to address the case in the third row of Figure 10. In this case, when the limiter is removed, the system of equations has no solution; one can see this as the limiting case of $\zeta_{\text{max}} \to \infty$, in which $u_* \to 0$, but $u_* = 0$ cannot be a solution since ζ (and hence, $\boldsymbol{f}$) is undefined at $u_* = 0$. When the limiter is applied, the system of equations is modified as the reviewer rightly notes. The limited system takes the form

$$\boldsymbol{x} = \boldsymbol{f}(\boldsymbol{x}; \boldsymbol{\xi}, \zeta_{\text{max}}) \tag{1}$$

and appears to have a unique solution for each value of ζ_{max}, as shown in Figure 10. From a purely mathematical perspective, this limited problem appears to have a unique solution. However, because the limiter is an artificial construct and has no physical meaning, this scenario is still not satisfactory.

We suspect, in part due to the strong correlation of low wind speed with the solution manifold $|\zeta| = \zeta_{\text{max}}$ (see Figure 12), that these issues arise in the free convection and extremely stable limits. However, we agree with the reviewer that this is not definitive evidence, and such evidence would be difficult to provide. We have thus greatly softened the language here (L514-520 in the revised manuscript).

- "In short, the proposed "adaptive" bound appears to mitigate the original $\zeta_{\text{max}} - 10$ "necessary evil", but not to eliminate it. Unless I missed it, the manuscript does not provide a quantitative assessment of how effective the adaptive procedure is. Figure 13 appears to be shown without the adaptive procedure: is that correct? A version of the same figure with the adaptive procedure, or a before/after comparison showing how often the fallback $\zeta_{\text{max}} = 10$ is invoked, would be very informative."

  Response: Thank you for pointing this out. Figure 13 is shown without the adaptive procedure since it is using the default algorithm in Alg. 1 with a larger number of iterations to diagnose how often the limiter is reached in the CTRL simulation. We have added a table, Table 1, which shows a count of the number of calls to the turbulent flux algorithm with oscillations or $|\zeta| = \zeta_{\text{max}}$ for the original algorithm (Alg. 1) and the revised algorithm (Alg. 3). The number of fallback cases is reduced by almost 90% with the revised scheme.

5. Minor comments:

- “Line 93 sets the tone well, but in my view, is too verbose in some respects and not explicit enough in others...”

  Response: We have removed the definitions of *uniqueness* and *existence* since we agree that they are self-explanatory. The definition of *equations underlying the turbulent flux parameterization* has also been removed and folded into the definition of the *turbulent flux parameterization*. We have also expanded on the definition of *well-posedness* to include continuous dependence on the problem data.

- “Line 106: does this have implications for the turbulent fluxes themselves? In other words, would the same turbulent-flux parametrisation/algorithm be used in a coupled run? Please make this explicit.”

  Response: The same turbulent flux parameterization and algorithm would be used in a coupled run. The only difference is in the inputs to the algorithm. In the uncoupled case, sea surface temperatures come from observation-based forcing data. In the coupled case, they come from the ocean model. This has been made explicit in the revised manuscript (L129).

- “L. 130: turbulent flux algorithms determine C_D, C_H, and C_E so that Eqs. (1)-(3) can be expressed in terms of large-scale prognostic variables.”

  Response: We have utilized this better wording in the manuscript now.

- “Line 135: terms such as “interdependency” or “cyclic dependence” may be useful here.”

  Response: Thank you for the suggestion. We’ve incorporated the term “cyclic dependence” in the revised manuscript.

- “L. 146: the relevant reference for (11) is (Obukhov, 1971).”

  Response: We have included this citation in the revised manuscript.

- “L. 164: which is discontinuous with respect to ζ.”

  Response: We have added this detail in the revised manuscript.

- “L. 179: that’s also the case in virtually all bulk algorithms, I believe.”

  Response: We have made note of this in the revised manuscript (L205).

- “Lines 181–183: what does the reference to Gauss–Seidel procedures add here? More importantly, the manuscript says little about how the iterative methods are initialised, even though for nonlinear problems the initial guess can be crucial for convergence.”

  Response: We have added some more details to the discussion of the Gauss-Seidel procedure (L207). In summary, we have provided some background on sufficient conditions for the fixed point iteration (e.g. contraction mapping) to converge to a unique solution. While the E3SMv2 procedure at first glance appears to use a fixed point iteration, the sequential application of independent fixed point iterations is actually a Gauss-Seidel method whose sufficient conditions for convergence to a unique solution are slightly different. While these details do not have major implications in the current work, we feel it is important to be precise about the exact numerical methods used so that future analysis of these methods is easier.

The initial guess for the iterative method is derived from neutral stability conditions, as noted in L180 of the original manuscript. Explicit expressions for the initialization are provided in L2 of Alg. 1.

- "L. 191 – 192: this is a major flaw that is worth being more underlined."

  Response: We have added some discussion around this point in the revised manuscript (L222), namely that two iterations likely contain large approximation error and lack of convergence checking makes diagnosing issues like the ones highlighted in this work much more difficult.

- "L. 240: for every point satisfying $\Delta\theta = 0$. In practise, this only happens when $\Delta\theta = 0$, doesn't it?"

  Response: In theory, this happens when $\Delta\theta = 0$. In practice, the numerical issues appear to happen when $\Delta\theta$ is close to 0 (see Figure 3), which we surmise is due to fixed point iterates jumping back and forth between the two branches of $\hat{C}_{\mathrm{HN}}$. We have mentioned this in the revised manuscript at L323-327.

- "Line 245: do $\boldsymbol{x}_i$ refer to $(U_{10N}, u_*, \theta_*, q_*)$, and $\boldsymbol{f}_i$ to the four sub-equations in (13)? Please make this explicit for the reader. Also, please make typography more consistent: $\boldsymbol{x}$ and $\boldsymbol{f}$ for the four-component vector (in bold), x_i and f_i for their scalar components (in normal font)."

  Response: We have made notation more clear and consistent so that $\boldsymbol{x}$ and $\boldsymbol{f}$ denote the four-component vectors and x_i and f_i denote their scalar components. The original manuscript had already defined $\boldsymbol{x}$ and $\boldsymbol{f}$ explicitly, and we have added some description in the same location (top of Section 2.4) defining x_i and f_i.

- "L. 250: near neutral, isn't it?"

  Response: Yes, we have amended this in the revised manuscript.

- "Figures 3, 7 and 12: hard to understand (1) There are too many points for the scatter plot to be informative. I think having the proportion of oscillatory condition per (x, y) bins would be much more relevant. (2) Histograms are not clear to me. I would put them on a separate one-row figure, with explicit x-labels and one shared [0; 1] y-label"

  Response: We have simplified Figures 3, 7, and 12 so that they now only contain histograms of U, $\Delta\theta$, and Δq in a single row. The bars now represent percentage of calls to the turbulent flux algorithm in each group which have values in each bin. The captions of these figures have been modified accordingly. We have removed the pairwise scatter plots since, as you noted, they have too many points to be informative and the original manuscript did not contain any textual references to them.

- "L. 275 – 277: this is a bit hand-wavy to me. Yes, the main non-converging points are in that $\Delta\theta$ band, which wouldn't it be possible to directly draw a map of non-convergence via directly binning with respect to geographical position (lat/lon)?"

Response: We have replaced Figure 4 with a new figure showing percentage of calls to the turbulent flux algorithm in each location experiencing an oscillatory iteration. We have also modified L275-277 to account for this change and note that non-convergence can and does occur outside of this $\Delta\theta$ range.

- "Lines 349–351: even if the algorithm converged very accurately, that would not by itself imply that the resulting turbulent fluxes are physically accurate; their parametrisation remains highly uncertain. This does not undermine the main point of the paper, but it is an important piece of context that could be stated explicitly. Of course, that still does not make it acceptable to hard code the iteration count to two!"

  Response: We have made this point more explicit by adding a sentence after Line 341 in the original manuscript noting that we are only discussing accuracy of the numerical methods in approximation solutions to Eq. (13), not physical accuracy of the turbulent fluxes (L390-392 of the revised manuscript).

- "Figure 8: the right panel would be easier to read if the x-axis showed average wall time per bulk-closure call rather than the total over the tested dataset. Arguably, this is only a relabelling change; the plot itself would remain the same"

  Response: We agree and have made this change to the x-axis.

- "Figure 13: could the lower end of the colour scale be given more resolution, perhaps with a nonlinear colour map? At present the lowest colour already corresponds to 10%, which is not especially low. Extending the scale down to about 1% would, in my view, be more informative. And are those actual percentages, or fractions with respect to 1?"

  Response: We have revised this figure so that the color scale follows a power law with the lower end of the range being allocated a larger portion of the total range. Note that the ticks on the color bar are still linearly spaced.