Transporting CRM Variance in a Multiscale Modeling Framework

Hannah and Pressel

We appreciate the comments from the reviewers as they have made some important points that have been addressed in the revised manuscript.

Response to Reviewer 1

In this article, https://agupubs.onlinelibrary.wiley.com/doi/10.1029/2021MS002892 , we have addressed a similar issue in superparameterized climate models, specifically for the transport of cumulus clouds in a high-resolution regionally super-parameterized model. We see the same issue of small-scale structures, in particular clouds, getting stuck in the CRMs where they are formed. We proposed a somewhat different solution to the one presented here, namely adjusting the total humidity variance (keeping the mean constant), in order to match the GCM liquid water content. Please consider referring to this article. I believe the similarities and differences in approach are interesting.

In our SP model with humidity variance adjustment, we saw that clouds and variability added through the coupling mechanism doesn't always persist in the CRM but can dissipate rapidly, if the receiving CRM is in a different state for example lacking the convection to support the clouds. It would be interesting to see if or how much this happens in the E3SM-MMF.

We are not quite sure how best to check for this behavior in our model, but we agree that this is an important thing to investigate. We do not have the infrastructure in place to output data at each CRM time step, otherwise we could easily check for a systematic dissipation of cloud within the CRM. We are especially curious if the additional consideration of static energy variance and momentum variance would help our implementation avoid the dissipation. Or perhaps our alternative “filtered” approach could help constrain this problem. This is something we will plan to look into once we have the infrastructure in place to do so.

Section 3.2 E3SM-MMF Description: Please mention whether the CRM uses the same vertical grid as the global model or something different.

This is a somewhat confusing aspect of our MMF and its predecessors, so we often avoid discussing it. The GCM uses a hybrid vertical coordinate that transitions from pure terrain following at low levels to pure pressure at upper levels. The CRM (SAM) is anelastic and uses a height vertical coordinate. Ever since the first incarnation of SP-CAM there has been no special consideration for this difference in vertical coordinates. Each time the CRM is called, the altitude and background pressure and density profiles are taken from the current GCM state and “updated” without any adjustment. In this sense the CRM always shares the GCM levels, despite having a different vertical coordinate.

At one point we decided to take a closer look at how this sudden update of height, pressure, and density affected the CRM solution. To investigate this we simply observed how the CRM reacted in the instant after the update by creating special output files with a snapshot of the CRM state at each CRM time step. To our surprise we couldn't identify any sort of shock to the system. While watching animations of the CRM as it evolved through this update step our eyes could not even identify where the update happened. Some have commented that this is not surprising given how the anelastic system works. So it seems that this quirky approach is harmless.

We have added the following details about the vertical coordinate to Section 3.2:

> *The global model uses a hybrid vertical coordinate that transitions from terrain following near the surface to a pure pressure coordinate near the top. The CRM employs the anelastic approximation and uses a height vertical coordinate. The mismatch in vertical coordinates turns out to be inconsequential for the MMF coupling scheme because the CRM background pressure and density can be updated with profiles from the GCM at each coupling step without causing an unphysical shock to the system.*

3.4 Satellite Observation Data: When re-gridding 1x1 degree satellite data to ne30pg2 of comparable grid size, isn't there a risk that the re-gridding operation may affect the presence of checker-board features? Regridding the much finer IMERG data to ne30pg seems safer.

This concern is addressed in the previous paper on our checkerboard detection method. Figure 6 of the previous paper is reproduced below. Comparing the black and gray bars provides a way to characterize the impact of regridding.

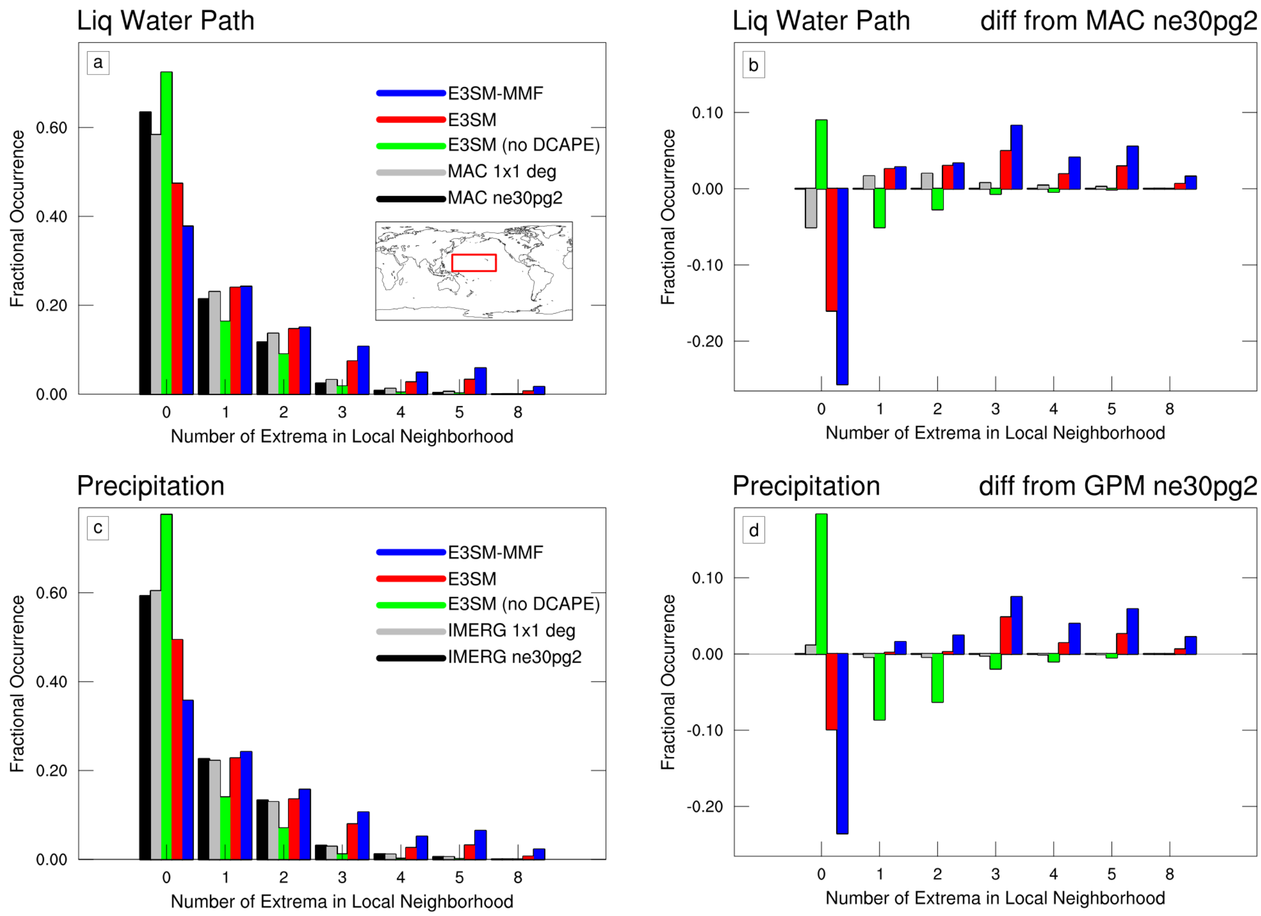

As an aside, another thing we explored but did not include in our previous paper was the question of how a pure checkerboard signal would be affected by the remapping. To explore this question we generated a pure checkerboard signal on the same 1-degree grid used by the MAC liquid water path data as shown in the top panel of the figure below. We then remapped it to the model's ne30pg2 grid (middle panel) and applied our checkerboard detection algorithm (bottom panel). If the checkerboard signal were able to be preserved by the remapping, then the bottom panel would be 1 everywhere, but clearly the checkerboard signal is mostly lost due to the remapping. However, we don't feel that this result affected our conclusions, so it was omitted.

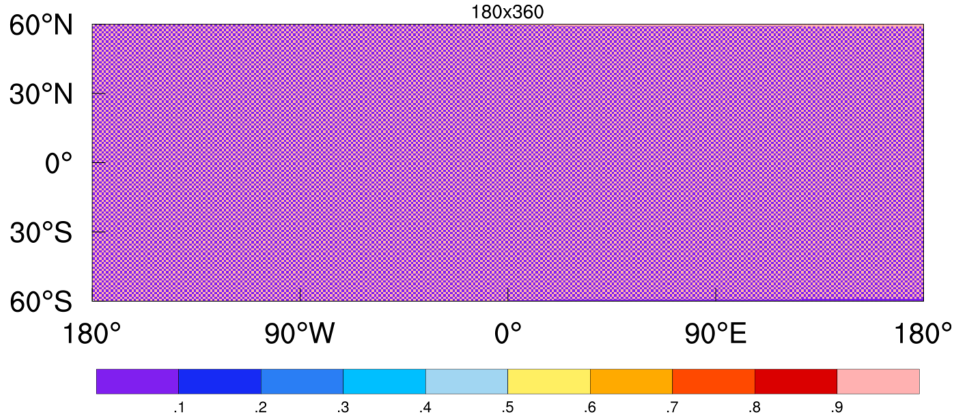

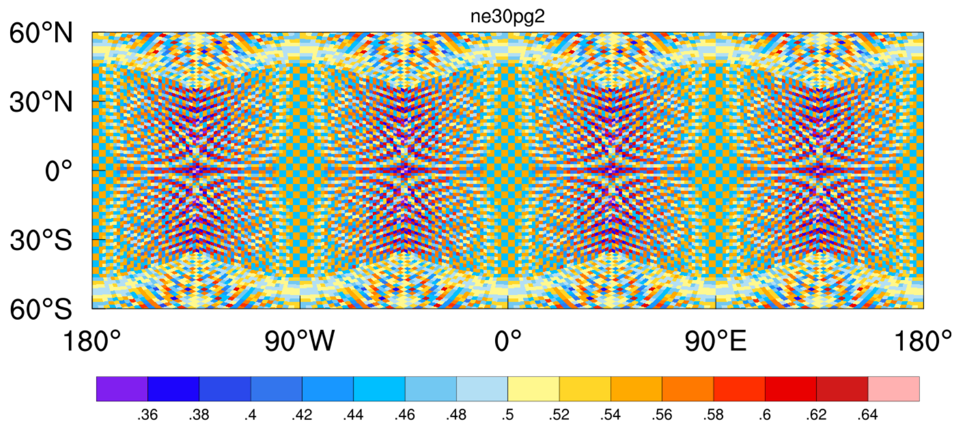

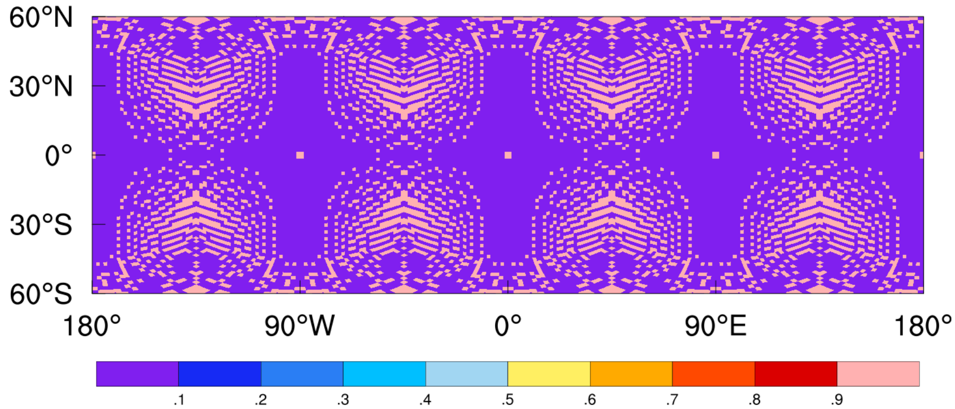

We performed a few other tests like this, but ultimately we decided that the analysis was not insightful and, thus, did not discuss it in the paper. We also experimented with running our detection algorithm on randomly generated data to provide a sort of baseline, but this also did not provide any additional insight.

We have added a note in the text explaining that the impact of regridding was explored previously and was determined to not qualitatively affect the results.

Note that the act of data regridding does affect the results of our checkerboard detection method. This was explored in more detail by Hannah et al. (2022) and the results showed a small impact that did not affect the conclusions. It is important to put the observational data on the model grid in this case in order to determine the level of naturally occurring noise that should be expected at the scale of the model grid.

Line 9: The first sentence is hard to parse because one can read it as "... cannot otherwise resolve convective scale circulations by coupling ..."

Good point. We have reworded this sentence.

Eq. (2) B_g should be B_G, or is not defined

Corrected.

line 65: <q_C^n>^n one n too much

Corrected.

Eq. (4) Left-hand term has a q without index, should be q_G?

Corrected.

Eq. (11) What is F?

Corrected.

line 93: "Since the CRM columns do have any specific location within the parent GCM column," do -> do not

Corrected.

Line 180: Unclear "gains performance by limiting radiation to operate of even subsets of the CRM domain"

This sentence has been rewritten into its own paragraph as follows:

"*E3SM-MMF uses a simple technique to reduce the number of radiative transfer calculations in order to boost throughput. Rather than performing radiation calculations on each CRM column separately, these calculations are performed on the average state of four groups of sixteen adjacent CRM columns and the resulting radiative tendencies are applied homogeneously back to the same group (Hannah et al., 2020). This provides a significant boost to the throughput without qualitatively affecting the model solution (not shown). A forthcoming paper will provide a* *detailed exploration of the model sensitivity to how these radiative groups are configured.* ”

Figure 4. Is the data averaged in time before or after checkerboard detection?

The checkerboard detection is always done on daily mean data. This has been clarified in the text.

Figure 9. Axis labels would be helpful, and a mention in the caption of what wave numbers are considered (e.g. which direction).

Axis labels have been added and the following sentence has been added to the caption to clarify that positive wavenumbers correspond to eastward propagation.

> *“Positive (negative) wavenumbers indicate eastward (westward) propagation.”*

Line 132: "However, experiments with this approach exhibited some odd behavior that we could not fully explain" and Line 140: "with some subtle, but unique, changes to certain climatological features." I'd recommend either saying more or less about these. Now I wonder what exactly happened in these two cases, while also suspecting that it may not be particularly relevant.

Looking back at this, I can’t remember what features I was referring to, but the more important point was that this “spectral variance transport” was simply just not effective at addressing the checkerboard patterns as the bulk or filtered method. We have rewritten this to highlight the more important result that this method was simply not effective.

Response to Reviewer 2

My first general comment concerns the stand-alone aspect of the submission. The presentation follows already published manuscripts concerning E3SM-MMF and reports developments motivated by those publications, up to the point of repeating some of the figures. I feel the authors assume that readers are familiar with the past work, and thus they use terms and concepts that are unclear unless you are familiar with those past manuscripts. Since I was not familiar, I had to go back to those previous manuscripts. I feel this needs to change to make the paper more stand-alone. I have several specific questions and suggestions below that should improve that aspect of the manuscript.

The dependence of this manuscript on the previous one was intentional since they are both motivated by the same problem, but the reviewer makes a good point that we can do better to explain certain things in more detail to avoid the reader needing to read the previous paper unless they are interested in some of the detailed results. We have written a few areas of the text with this in mind.

My second comment concerns conservative properties of the model with variance transport. The original GCM-CRM coupling conserves energy (in its simplified representation) and water substance by design. Is that true for the system with variance transport? I think it is by design as well, but perhaps worth commenting on it and maybe even illustrating by some additional analysis.

The reviewer is correct in their intuition that the variance transport does not affect energy or water conservation. The variance coupling does not affect the CRM domain mean and therefore does not affect conservation at that scale. The local conservation within the CRM is less straightforward, but the conclusion is similar. The main prognostic variables of the CRM, liquid-ice static energy and total water, are appropriately conserved within the CRM, so the local sources and sinks created by the variance injection, with no effect on the domain mean, do not affect the general conservation of these CRM fields. We have added the following sentence to section 2.2 to address this concern:

> "*Note that the variance transport does not affect the mean state of the CRM, and, thus, does not affect energy or mass conservation.*"

1. Title: I suggest replacing "CRM" with "cloud-scale" or "convective-scale" to avoid the acronym that might be unclear to some readers.

The original draft title actually did use one of these terms. We went back and forth about this because many LES modelers apparently hate the term "cloud resolving" and don't consider the convective circulations "real clouds", or something along those lines (I've never understood this complaint). In any case, we determined that we should not use "cloud", and "convective" seemed too ambiguous, so I settled on just using "CRM" since it is such a common term that people associate with a certain range of grid spacing, despite the fact that some people still dislike it. So for these reasons we prefer to keep the original title.

2. A general comment on the introduction: I do not think the key problem with MMF is the presence of the spectral gap between scales resolved by GCM and those in CRM. The key is that the CRM domain is periodic and the cloud-scale signal cannot propagate from one GCM box (i.e., one CRM) to the CRM in the next GCM box. In other words, small-scale perturbations are trapped in a GCM. This is why transporting variance helps, correct? One way to stress that in the discussion would be to emphasize that CRMs feature periodic lateral boundaries. For instance, adding cyclic or periodic in line 11 would help. Periodicity is important from the energy conservation point of view. Making CRMs open brings essential problems with keeping track of what comes in and comes out. The sentence in line 24, "Another consequence…" is simply incorrect. Lack of advection has little to do with the scale gap; it comes from periodicity of CRM domains. Another comment is that it would help if the checkerboard pattern is already documented in the introduction to set the readers on the right pass. This is only mentioned now (line 19-23), but adding a figure from a previous paper would help. And, again, I do not agree that the key is the scale gap, the lack of large-scale advection is the culprit.

We feel that this is really a semantic issue with no right answer, but we disagree with the reviewer in that we feel the CRM's periodicity is not a problem in itself. If we relaxed the need for periodicity we would need a completely different paradigm to couple the two models, as was explored in the "Quasi-3D" framework invented by Akio Arakawa, so the issue goes far beyond the horizontal boundary conditions of the CRM. We prefer to think of the lack of large-scale advection of small-scale features in the CRM as a product of the scale-gap, and we do not feel the reviewer has made a strong case for why our statements are "simply incorrect".

Ultimately, this is about how to properly frame the problem, so perhaps focusing on the need for large-scale advection of small scale signals is more important than pointing out the scale gap, but we still feel that the scale gap is a useful conceptual framework.

In regards to the notion that the periodic boundary conditions in the CRM are problematic, we ran an experiment before stumbling onto the variance transport concept that tried to effectively relax the need for periodic boundaries. Our approach was certainly ad-hoc, but seemed to work for the question we were asking at the time. The idea was to enlarge the CRM domain and only explicitly couple a portion of it while relaxing the rest of the domain back to some state (not exactly clear what the target should be for that). Technically we still had periodic boundaries, but the relaxation effectively made it so that signals leaving or entering the coupled subdomain were not being recycled within the domain, and hence the "variance trapping" could not occur. This approach is computationally problematic since it requires a larger CRM domain, but the basic idea seems to be effective at relaxing the requirement of periodicity, and allows variance to leave the domain rather than being trapped. However, in either case the small-scale fluctuations still cannot be directly advected by the large-scale flow, so the scale gap is still problematic without periodicity.

We have added a sentence to the second paragraph to expand on the checkerboard noise pattern, which includes a reference to the new inset of Figure 1b (see below).

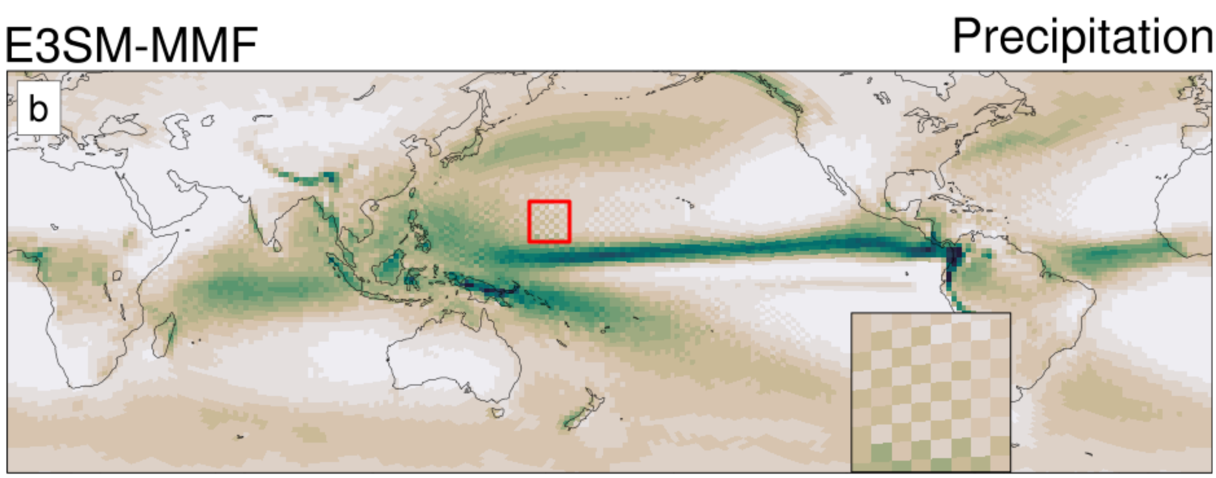

3. Section 2.2. I think it would help if the technique is illustrated by a figure. If I understand the approach, you simply "scale up" the variance in each CRM based on GCM advection from the *neighboring columns, correct? The key is that alpha is advected by the GCM, correct? If the CRM field (at a given height?) is q, you make it alpha q, correct? Can this be shown in a figure? Also, alpha is height-dependent, correct? It would help to state this clearly.*

The reviewer is correct that alpha is a scaling parameter, but the GCM advects the variance field, rather than alpha. The alpha values are calculated at each CRM time step in order to scale the existing anomalies. Another way to think about the derivation of alpha is to ask "how much do we need to scale the existing perturbations to get the change in variance dictated by the variance forcing (F). We have attempted to clarify this in the revised text. The reviewer is also correct that alpha is height-dependent, we have noted this in the revised text.

It is unclear how we would clearly illustrate the variance transport scheme. Our worry is that a collection of arrows and symbols would require many simplifications that may lead to confusion. Thus, we feel that the description in the text is the best way to communicate the concept.

If we had the infrastructure in place to run the MMF in a regional (planar) configuration with idealized initial conditions and forcing we could probably create a good visualization of how variance transport affects the solution, similar to the figures in this paper:

> Jansson, F., van den Oord, G., Pelupessy, I., Chertova, M., Grönqvist, J. H., Siebesma, A. P., & Crommelin, D. (2022). Representing cloud mesoscale variability in superparameterized climate models. Journal of Advances in Modeling Earth Systems, 14, e2021MS002892. https://doi.org/10.1029/2021MS002892

But, unfortunately, we don't have such infrastructure in place.

4. What is F in (11)?

Thanks for pointing this out. This is the variance forcing term from the RHS of (8). We have edited the text to clarify this.

5. Section 3.1. I feel a figure describing the detection would help. This is not clear unless one goes back to Hannah et al (2022). A figure or two from that paper would help.

We feel that referring the reader to Hannah et al. (2022) is preferred so we can avoid duplicating material. We have rewritten parts of this section to provide more details on the detection algorithm and refer them to specific figures from our previous paper.

We should also note that we made an error in our original description of the checkerboard detection, which has been corrected in the revised manuscript. We said that the local neighborhood anomalies were calculated relative to the neighborhood mean, but in fact we calculate these anomalies relative to the central value.

6. Section 4.1. Several statements in this section are hardly evident in the figures. Line 234: perhaps a hint of a double ITCZ? L. 235: I do not see the checkerboard pattern. My suggestion is to improve the figure, maybe with a small insert, so the features mentioned are better *documented. I think Fig 2 does show that pattern, correct?*

The pattern should be evident in the differences shown in Figure 2, but we primarily intended this figure to highlight the systematic regional differences that result from enabling variance transport, so it's not quite as important for the reader to see the checkerboard patterns. However, seeing the checkerboard patterns in Figure 1 is still important. Thus, we have added an inset to each panel of Figure 1 (see new version below).

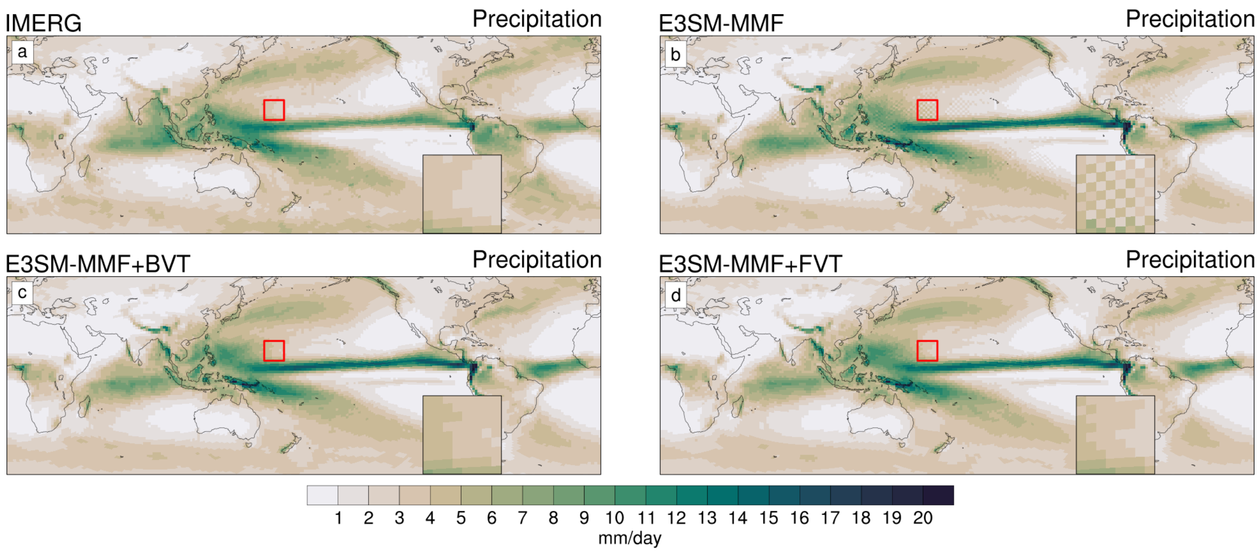

7. Section 4.2. How is "fractional occurrence" defined? It is unclear to me what Fig. 3 shows. How the "extremum in the local neighborhood" is defined (in IMERG and in the model results)? Figure 4 shows clear differences, but if I do not know what the fractional occurrence is how can I gauge the significance of those differences? Fig. 5 provides a clear impact of the improvement and good comparison with observations.

Fractional occurrence is simply the number of observations of a pattern (or set of patterns) normalized by the total number of valid observations at each location or across the region of interest. Identifying "valid" observations is especially important when dealing with the GPM data since there are many instances of missing data. We have added a better explanation of "fractional occurrence" to section 4.2 in the revised manuscript.

It's unclear how we might estimate the statistical significance of the fractional occurrence differences. We do not have a good way to show an "expected" value of the fractional occurrence in Figure 3 or 4 aside from comparing to what we can detect in the satellite data.

I do not understand what Fig. 6 shows, please explain.

To provide clarity on Figure 6 we have written the beginning of the paragraph as follows:

> "*The use of daily mean data for pattern detection in the previous analysis was meant to facilitate the comparison with satellite data, but it is insightful to consider finer time scales in order to see how the checkerboard patterns "spin up" at the start of a*

simulation. To investigate this question we reran E3SM-MMF for 10 days with output at every time step (20 min) and ran each snapshot through our detection algorithm. Figure 6 shows a time series of the fractional occurrence of partial checkerboard patterns at each model time step from the beginning of each simulation. Data is restricted to ocean regions between 60S-60N and the horizontal black dashed lines indicate the climatological fractional occurrence of partial checkerboard patterns from the corresponding satellite data for reference. Interestingly, the checkerboard signals are detectable in the first few time steps for liquid water path, but are slower to spin up in the precipitation field. After the initial spin up the occurrence of partial checkerboard patterns continues to increase over the first 4-5 days before leveling off. The speed of the initial spin up is consistent with the authors' experience that convection within the CRM develops very quickly after initialization. This further illustrates that while the CRM variance transport method can ameliorate the checkerboard pattern persistence and yield a much smoother climate, the model still produces a relatively noisy solution on short time scales relative to satellite observations."

We have also added axis labels to Figure 6 and edited the caption.

*8. Please clearly define what each panel of Fig. 7 is showing. Also, do the numbers shown have some relevance to the real world? For instance, the maximum temperature variance (temperature or potential temperature?) is around 0.6 K**2. Does this mean that the CRM variance is less than 1 K? How this compares to either observations or cloud-scale simulations of tropical convection?*

This is something we glossed over since we are so familiar with this type of analysis, but we certainly agree with the reviewer that more discussion is needed. The "T" variance here is actually from the liquid-ice static energy, which we have clarified in the figure caption. The variance tracer values are calculated directly from the CRM fields, and then modified by transport and diffusion in the GCM. We do not know of a suitable dataset to make a direct comparison with observations. We could make a comparison to a global cloud resolving model, which could certainly provide an interesting and valuable insight, but we consider this outside the scope of the present study. We have added a sentence about this possibility for future work in this section.

Can you explain why GCM and CRM transport panels seem to have the same pattern but the opposite sign? What is the reason for that? Some physical interpretation would be good. This is briefly discussed, but I would like to see more physical interpretation. Similar comments apply to Fig. 8 (please define all panels in the figure caption).

The apparent cancellation of CRM forcing and feedback tendencies is an expected outcome of the design of the coupling scheme that ensures no drift between the CRM domain mean and its parent GCM column. We have edited the text to clarify this in the discussion of Figure 7 and updated the caption of both Figure 7 and 8.