General comments:

This paper reports on simulations, using WRF, of a case study in the Australian GBR region, of how cloud microphysical and macrophysical properties are affected by marine cloud brightening sea salt aerosol injections. Three injection cases are simulated, all with the same total mass of injected sea salt, but with different distributions of point sources (EXP20: 75 sources; EXP40: 12 sources; EXP100: 3 sources). The EXP20 case produces the largest cloud albedo changes, and this is shown to be mostly via the Twomey effect. The paper further breaks down the results from the EXP20 study to quantify differences in cloud brightening driven by three metrics for meteorological controls: above-cloud humidity, wind sheer below vs above cloud, and boundary layer estimated inversion strength (EIS).

In particular, the contribution of examining the effect on MCB efficacy of spray system spatial distribution is a valuable addition to the literature, and overall the analysis is technically sound. The paper is generally well-written and well structured, and the figures clear. However, I think the paper includes statements that I think are either too strong or unsubstantiated by the presented analysis, and in a few places where wording needs clarification.

Specific comments:

1. lines 21-24: It's not clear what's meant by "intensity". Is this just referring to the number of emission points? Or the rate of emission at each emission point? It's not clear from the text here that all simulations have the same total emissions, just distributed over a different number of emission points, as described in the main text. This text should be edited for better clarity.

> The "intensity" refers to the emission rate at each station. As the total emission is held constant across all simulations, the per-station emission rate increases when fewer stations are used (e.g., three stations in EXP100), and decreases when more stations are deployed (e.g., 75 stations in EXP20). Detailed values of the emission rate for each experiment are provided in Table 2.
>
> The text has been revised as follows:
>
> "A control simulation representing a non- to weakly-precipitating shallow trade-cumulus regime is compared with three MCB sensitivity experiments: a densely distributed (20 km apart), moderate emission flux scenario (EXP20), a sparsely distributed (100 km spacing), high emission flux scenario (EXP100), and an intermediate configuration (EXP40)."

2. lines 134-136: What size aerosol are emitted? This is important, given the strong dependence of aerosol activation on aerosol size.

> Thank you for this important comment. In the Thompson aerosol-aware microphysics scheme used in this study, the emitted aerosol is not represented as an explicit size-resolved aerosol distribution. Instead, aerosol perturbations are introduced as changes in the bulk number concentration of water-friendly aerosols, which are available for cloud droplet activation.
>
> As the scheme does not explicitly resolve aerosol size distributions, activation is not computed using a size-resolved Köhler framework. Rather, the number of activated cloud droplets is diagnosed using a bulk formulation, in which activation efficiency is implicitly controlled by aerosol number, vertical velocity, and thermodynamics.
>
> We acknowledge that this bulk treatment does not allow us to assess the sensitivity of activation to aerosol size, which may be important for real MCB particle design and deployment. This is noted as a limitation of the current modelling framework.

3. line 193-197: In reference to the comparison of the simulated (panels a and d) vs observed (panels b-c and e-f) cloud fields, shown in Figure 4: I'm struggling to agree with the authors that the simulations capture the cloud field "reasonably well". The cloud fraction in the study area is quite different (given later as ~7% in the simulations but what looks more like 50-60% in from Himawari obs), as is the morphology of the cloud fields (very small, broken, evenly distributed Cu in the simulation, vs broader and more variable areas of cloud cover in the obs).

This leads me to think that the boundary layer thermodynamics might be different in the model than in reality. Figure 3 shows good what looks like very agreement between the simulated and observed temperature, dew point and wind profiles, but for these clouds what will most matter is what's happening in the lowest few km of the atmosphere. In Figure 3a, the simulated and actual winds below ~800hPa are actually pretty different; in Figure 3b it's hard to see what the sounding winds are doing.

-> I'd suggest adding two additional panels to this figure, reproducing the current panels a) and b) but zoomed in on, e.g., 700mb to the surface.

-> Further, I think this statement needs to be softened. The truth is that it is really challenging to reproduce the specifics of a given cloud field, even with a perfectly initialized simulation. This can be acknowledged in the context of not overstating how well the model is reproducing reality.

> Thank you for this insightful comment. We agree that reproducing the detailed cloud fraction and morphology of a specific cloud field is inherently challenging, particularly for shallow cumulus regimes where cloud organization is highly sensitive to small-scale variability and stochastic processes.
>
> In response to the reviewer's suggestion, we have added zoomed-in panels of the sounding comparison (600 hPa to surface) in Figure 3. These additional panels show that the model reproduces the lower-tropospheric thermodynamic structure (temperature and moisture) reasonably well, while some differences remain in the wind profiles around 800 hPa.
>
> We acknowledge that the simulated cloud fraction is lower and the cloud field appears more homogeneous compared to Himawari-8 observations. Therefore, we have revised the manuscript to soften the statement regarding model performance as follows. Rather than focusing on detailed cloud morphology, we emphasize that the model captures the key thermodynamic environment and the presence of shallow cumulus conditions, which are the primary controls for aerosol–cloud interaction processes examined in this study.
>
> "Overall, the model underestimates the cloud fraction and does not fully reproduce the observed spatial organization when compared with Himawari-8 observations. In particular, much of the high-level cloud cover within the domain is underestimated or missing, leading to discrepancies in the representation of fine-scale cloud structures. Such differences are expected given the inherent difficulty in reproducing the exact realization of convective cloud fields. However, the model captures the key characteristics of the maritime shallow cloud regime visible in the true-color imagery and brightness temperature proxy, which is the primary focus of this study. In particular, the lower-tropospheric thermodynamic structure (Figure 3c and 3d) is reasonably well represented, providing a physically consistent boundary layer environment for shallow cloud formation. Therefore, although the detailed cloud morphology is not fully reproduced, the simulation is considered suitable for investigating aerosol–cloud interaction processes under shallow cumulus conditions."

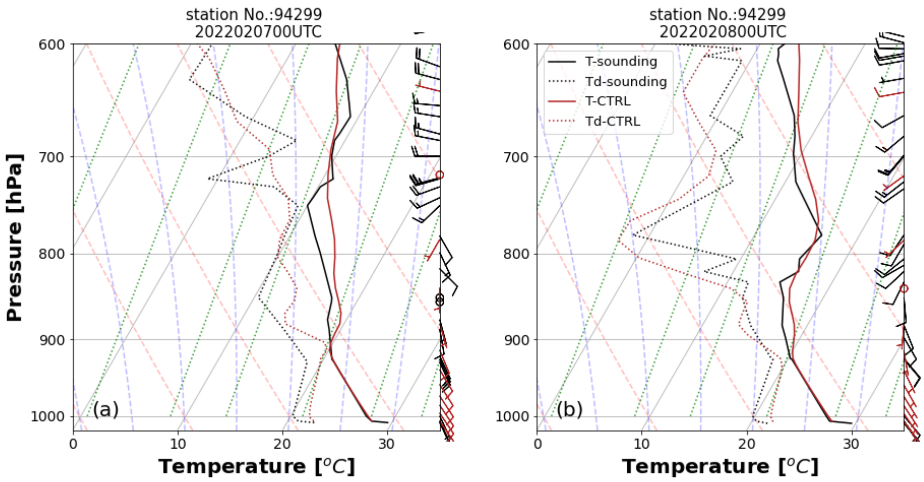

Figure R1: (a) Zoom-in comparison (surface to 600hPa) of the observed sounding profiles (black lines) from Willis Island alongside 1km WRF atmospheric profile simulations (red lines) from the nearest grid point on 07 Feb 2022 at 0000UTC. (b) same as (a), but for 08 Feb 2022 at 0000UTC.

4. Figure 5 vs Figure 6: Comparing these two figures it appears the model is only activating about 15% of the aerosol to CCN – an activation rate that seems very low. This could be due to the resolution of the model (1km in the inner domain; 5km in the outer domain), as the lower resolution could be leading to a low bias in updraft rates. Ultimately this will produce a bias in the mass of aerosol that needs to be emitted to achieve a given forcing. Perhaps in the discussion, the simulated aerosol activation rate should be put in the context of expected values for typical low marine stratocumulus and this potential source of bias acknowledged.

Thank you for this insightful comment. We agree that the initially reported activation rate (~15%) appeared low compared to expected values.

In the original analysis, aerosol concentrations were sampled at the surface level, while CDNC was derived from vertically integrated cloudy pixels. This inconsistency led to an underestimation of the activation fraction, as surface aerosol does not represent the population of particles available for activation at cloud base.

In the revised analysis, we instead diagnose aerosol number concentration in the sub-cloud layer (100–200 m below cloud base) and evaluate CDNC at cloud base (or the first in-cloud level). This approach is more physically consistent with activation theory, as it captures the aerosol population that directly participates in cloud droplet formation.

Using this updated methodology, the estimate activation fraction is approximately 30-35%, as shown in the Figure R2. This range is consistent with previous studies of marine shallow cloud regimes (e.g., McFiggans et al., 2006), particularly under moderately polluted conditions (e.g. MCB experiments) where activation fractions can exceed 30% depending on updraft strength and aerosol size distribution (McFiggans et al., 2006).

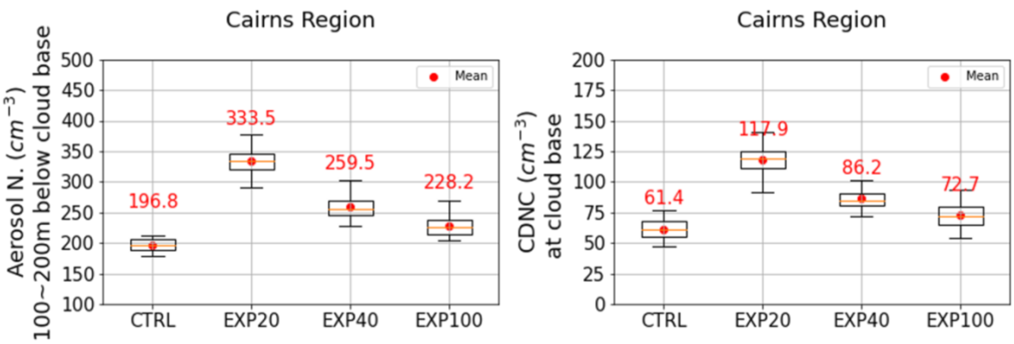

Figure R2: (Left) Aerosol distribution below cloud base from WRF experiments (Right) Cloud droplet number concentration at cloud base from WRF experiments. Time period for analysis is 08 UTC 07 Feb 2022 to 08UTC 08 Feb 2022. Red numbers are mean values of each experiment.

5. Figure 6 a-c and Figure 7a: Why show the *surface* aerosol concentration, rather than the concentration at cloud base? The variability in concentration at the surface isn't what matters; it's the variation in concentration at cloud base that will affect CDNC.

(Note that I'm asking this based on the assumption that "surface" means the value in the model's lowest layer, but maybe that's incorrect? If you're going to stick to showing 'surface' concentrations it would be good to at least define what you mean by 'surface'.)

Thank you for this helpful comment. We originally showed aerosol concentrations at the surface to illustrate the spatial distribution and configuration of the aerosol injection in each experiment. Here, "surface" refers to the aerosol concentration in the lowest model level.

We agree that aerosol concentrations near cloud base provides a more physically relevant representation of the aerosol field interacting with clouds. We have now revised Figures 7a to present aerosol distributions at cloud base (or the model level closest to cloud base, see Figure R3). Figure 6 keeps as initial to show the aerosol distribution at surface level. The corresponding text has been updated accordingly to reflect this change.

6. lines 241-244: The explanation given here of the precip scavenging efficiency at higher aerosol concentrations driving the difference in domain-mean concentrations across the different cases intuitively makes sense -- except that it is noted earlier that these are "non- to weakly-precipitating" clouds. Later, Figure S3 shows that the clouds were indeed only precipitating for a few hours in the middle of the simulation, and at very low rates (<0.01mm/hr). Can the simulated rain rate be used to quantify (even approximately) how much precip scavenging could be accounting for the differences in domain-mean aerosol concentrations across the different cases? And does this align with the simulated differences in concentrations? I think such an analysis should be included, especially given it affects the first main conclusion of the paper (lines 429-430)

More generally, is there a reason for ruling out other causes of the differences? e.g. Could coagulation or other processes be driving some of the difference? (This relates to the question of the injected aerosol size).

If a difference in precip scavenging is truly the main reason why the domain-mean aerosol concentrations differ, then the results shown here are very specific to the precip rates in this case, which brings into question whether the analysis of a single case allows one to reach a conclusion about the most efficient deployment configuration in general.

We agree with the reviewers that our previous interpretation overemphasised precipitation scavenging. Given the non- to weakly precipitating nature of the simulated case, wet scavenging by precipitation is unlikely to be the dominant mechanism explaining the lower domain-mean aerosol enhancement in EXP100. We have therefore revised the manuscript to avoid attributing this behaviour primarily to precipitation scavenging efficiency.

Instead, we interpret the results as arising from enhanced aerosol processing within concentrated emission plumes in the sparse-source configuration, which may lead to more efficient domain-wide dispersion and aerosol loss. Multiple processes may contribute to this behaviour, including cloud processing, dilution, possible coagulation, turbulent mixing, and deposition. In particular, processes such as coagulation may reduce aerosol number concentration without necessarily reducing aerosol mass, which could partially explain the differences between experiments.

However, the current model output does not fully resolve or diagnose the relative contributions of wet scavenging, dry deposition, coagulation, and transport. This limitation has now been explicitly acknowledged in the revised manuscript.

Lines 239-244 have now been revised to

"This behaviour arises from the nonlinear response of aerosol processing and removal pathways to local aerosol number concentration in the WRF Thompson aerosol-aware microphysics scheme (Wang et al., 2010; Weston et al., 2022). In the sparse-source configuration (EXP100), aerosol emissions are concentrated into fewer locations, leading to enhanced aerosol processing within these concentrated plumes. Multiple processes may contribute to this behaviour, including cloud processing, dilution, possible coagulation, and deposition. Given the non- weakly- precipitating nature of the simulated case, coagulation within the highly concentrated aerosol plumes is likely an important, and potentially dominant, contributor, as it can efficiently reduce aerosol number concentrations near the source regions and thereby limit the enhancement of domain-mean aerosol concentrations. However, it should be noted that the current model output does not fully resolve or diagnose the relative contributions of wet scavenging, dry deposition, coagulation, and transport."

7. lines 321-323: I have to object to the assertion that the cloud albedo exhibited "pronounced increases" in the injection cases relative to the CTRL case. The increases are well within the range of variability in the CTRL case (Figure 8e); in Figure 7e, only the EXP20 case has a cloud albedo that is notably and consistently higher with aerosol injection than without.

We thank the reviewer for this important clarification. We agree that the term "pronounced increases" may overstate the magnitude and robustness of the cloud albedo response across all experiments. It should also be noted that the statistics shown in Figure 8 have now been updated following the revised cloudy-mask methodology, in which cloudy conditions are defined using cloud optical depth (COD) > 1, as requested by other reviewers. The updated analysis still indicates relatively modest albedo responses compared with the CDNC changes, thus the overall conclusions remain consistent.

Accordingly, we have revised the wording in the manuscript to describe the cloud albedo response more cautiously and quantitatively.

We would also like to clarify that, although the domain-mean cloud albedo response is relatively modest, the shortwave cloud radiative effect (CRE) analysis introduced in the revised manuscript demonstrates more substantial radiative perturbations associated with aerosol injection. This difference partly arises because the synthetic cloud albedo in this study is diagnostically derived from cloud optical properties that primarily depend on cloud water path (CWP) and cloud droplet number concentration (CDNC), rather than being directly archived from the model output. In contrast, CRE provides a more direct measure of the net shortwave radiative impact associated with aerosol-induced cloud modifications, which is a direct output from the model. We retain the synthetic albedo analysis because CRE is only physically meaningful during local daytime, whereas synthetic albedo can also provide information on cloud optical properties during nighttime.

To avoid overstating the implications of the albedo analysis, we have revised the manuscript to focus primarily on the simulated CRE response and the process-level sensitivity of cloud radiative properties to aerosol perturbations, rather than emphasizing large albedo increases alone.

8. lines 431-432: This is an incomplete sentence.

Also, again I think that other than (marginally) in the EXP20 case, the cloud albedo increases were modest to (EXP100) negligible.

Thank you for pointing this out. The sentence has been revised as follow:

"Substantial increases in CDNC and cloud optical depth lead to brighter clouds and higher mean albedo under additional aerosol perturbations."

9. line 441, Conclusions: "The findings confirm that radiative forcing can be meaningfully enhanced under suitable conditions through spatially coherent aerosol seeding. As such, strategic deployment of MCB during favorable weather regimes holds promise as a targeted intervention to mitigate extreme heat exposure over sensitive marine ecosystems like the GBR".

I feel strongly that these statements cannot be made. No radiative forcing calculations are shown. The paper later asserts that the cloud albedo change from EXP20 vs CNTRL could produce a forcing "potentially on the order of several tens of W m^{-2}" but I don't see how this is the case. In EXP20, the cloud albedo increases to 0.51 from 0.48. If there was 100% cloud cover this would, ballpark, result in an increase in reflected solar flux of ~10 W/m2 (downwelling 340 W/m2*(0.51-0.48)=10.2 W/m2). But cloud cover isn't 100% -- it's, in this case, 7%.

Further: Being able to say forcing is *meaningfully* enhanced, and that MCB could actually mitigate extreme heat exposure in the GBR, would require showing that the resulting forcing would be sufficient to produce ocean cooling.

All of this is well beyond the scope of the paper – but so is the assertion.

Unless the authors want to add this analysis, they can only assert what the paper has actually shown, which is articulated nicely on lines 446-495, and is of sufficient scientific value to stand on its own.

We thank the reviewer for this thoughtful and constructive comment. We agree that our original wording may have overstated the broader climatic implications of the simulated radiative responses. In response, we have revised the conclusions to better align with the actual scope and findings of the study.

Following this comment, and in conjunction with the Comment 7, we have now added analysis of the shortwave cloud radiative effect (CRE) associated with the aerosol perturbation experiments. In addition, the cloud mask methodology has been updated following reviewer recommendations to define cloudy conditions using cloud optical depth (COD) > 1, which represents optically significant clouds and is commonly used to improve robustness in cloud detection and radiative analyses (Heidinger et al., 2020). The corresponding analyses in Figures 7 and 8 have been updated accordingly, although the overall statistical behaviour and conclusions remain largely unchanged.

As shown in the Figures R3 and R4 (updated Figures 7 and 8), the aerosol-seeding sensitivity experiments exhibit enhanced shortwave cloud radiative effects relative to the control simulation, indicating increased reflection of incoming solar radiation due to aerosol-induced cloud brightening. The largest CRE enhancement occurs during periods of high solar zenith angle near local midday, when incoming solar radiation is strongest. While the simulated case study covers only a limited daytime period, the results consistently demonstrate measurable increases in cloud reflectivity and shortwave radiative cooling associated with additional aerosols.

However, we agree with the reviewer that our simulations alone are insufficient to determine whether these radiative perturbations would be large enough to produce meaningful ocean cooling or mitigate marine heat stress over the GBR. Such conclusions would require substantially broader analyses, including coupled atmosphere–ocean modelling, longer temporal integrations, and regional energy budget assessments, all of which are beyond the scope of the present study.

Accordingly, we have revised the manuscript to avoid overstating the implications of the results. The conclusions now focus specifically on the demonstrated process-level sensitivity of cloud microphysical and radiative responses to aerosol seeding configuration and meteorological conditions, rather than asserting direct mitigation of extreme heat exposure over the GBR.

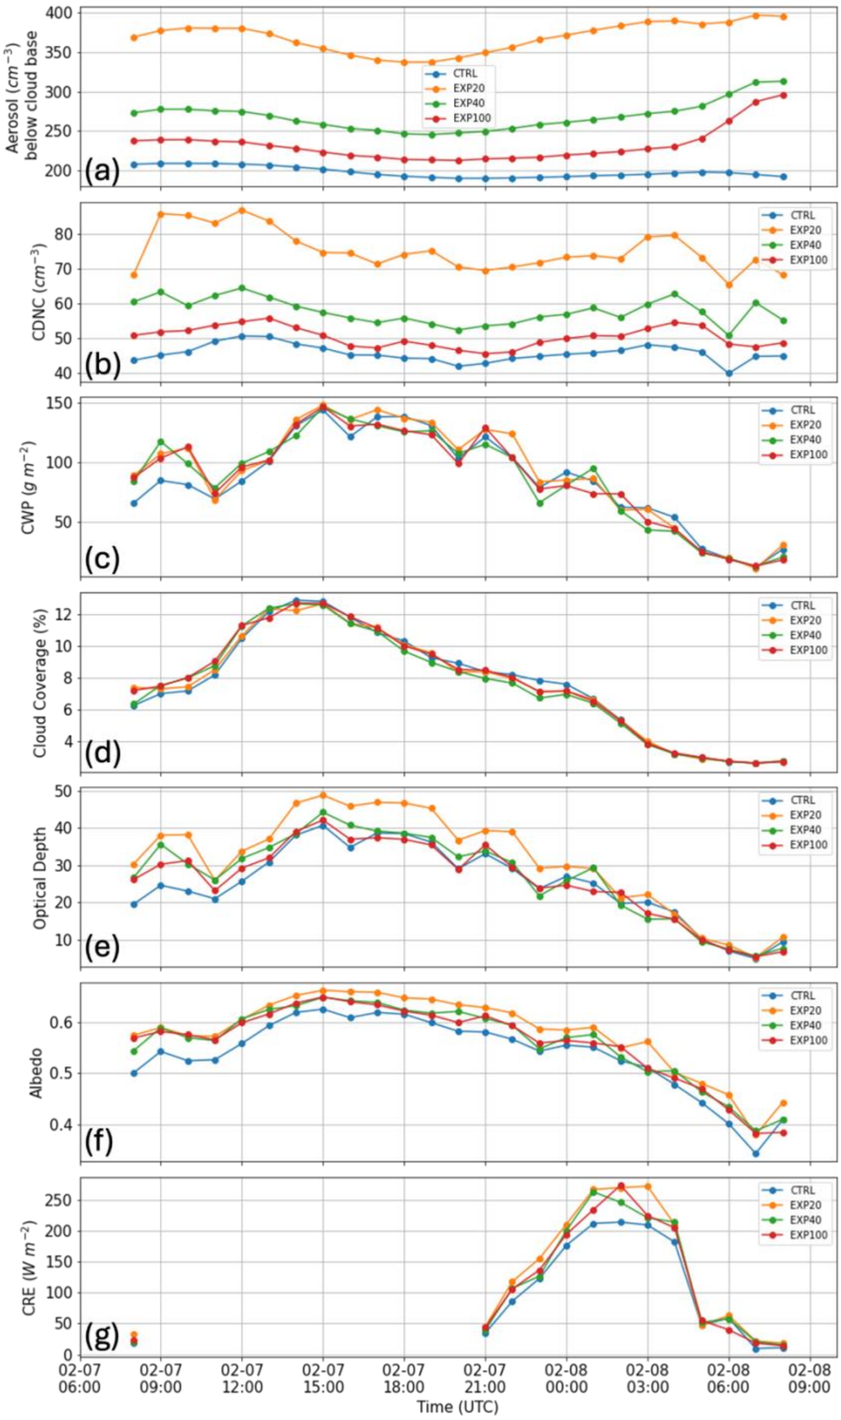

Figure R3: Time series of domain-averaged (a) aerosol number concentration below cloud base, (b) CDNC, (c) cloud water path, (d) Cloud Coverage, (e) Optical depth, (f) synthetic cloud albedo, and (g) cloud radiative effect over the Cairns Region. Note that all variables are averaged over cloudy grid points.

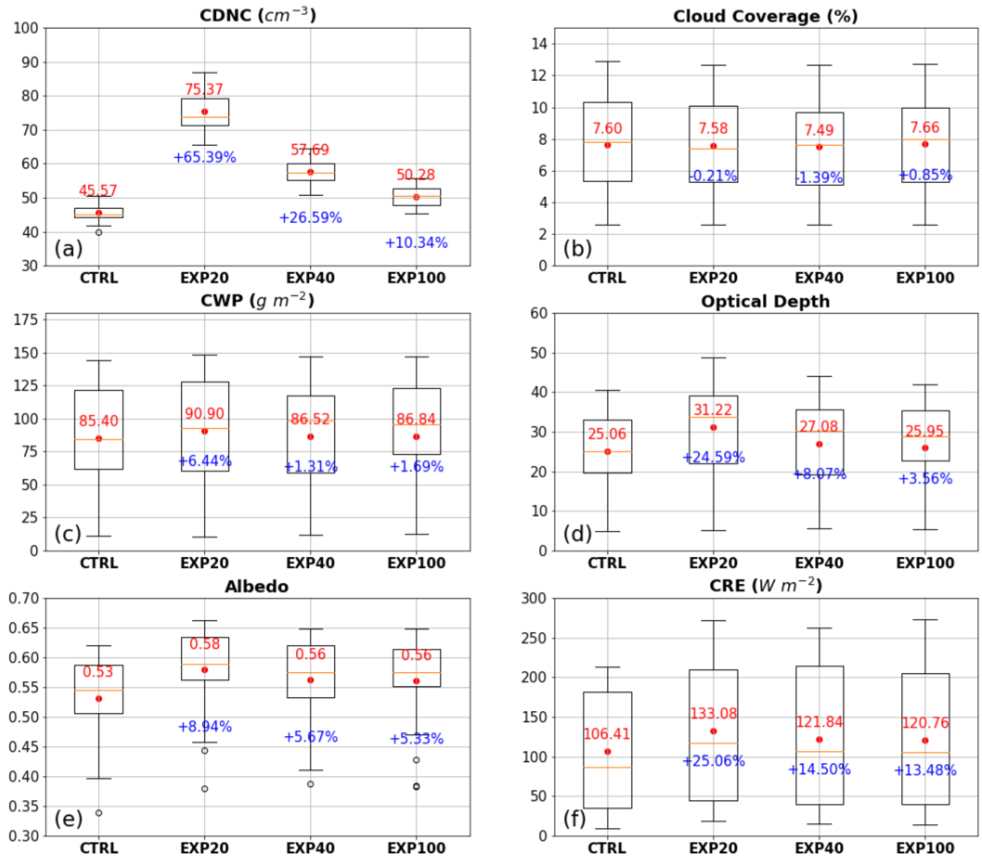

Figure R4: Boxplots of cloud properties over Cairns Region from CTRL and sensitivity experiments. Time period for analysis is 08UTC 07 Feb to 08 UTC 08 Feb 2022. Red numbers are mean values of each experiment, and blue numbers are differences with respect to CTRL experiment.

Technical comments:

- lines 68-69: A classic paper that should probably also be cited here is the Stevens and Feingold (2009; https://doi.org/10.1038/nature08281) paper that's also cited in another context later in the paper.

  The reference has been cited as suggested.

- line 184: skilfully -> skillfully

  Done.

- line 184-185: "The control simulation skilfully simulated the evolution of the synoptic-scale mean surface level pressure and surface winds (not shown)." The evidence showing this should be included in the Supplemental data/file.

As suggested, the relevant materials have now been added into Supplementary. The corresponding text has been revised too.

- Figure 5: It appears that none of the CDNC values in any of the panels exceed 150/cm3, so why not set the colorbar range from 0-150/cm3 (or maybe even 0-125/cm3)?

  The colorbar range has now been revised to 0-140/cm^3.

- line 231: "In the contrast, …" -> "In contrast, the …"

  Done.

- line 233: "exhibited" -> "exhibits" (for consistency present tense used elsewhere)

  It has now been corrected.

- line 435: "mid-level" is a bit ambiguous; I'd say "above-cloud"

  Thanks for your suggestion. It has been revised to "above-cloud".

- line 524: this is the first use of LWP; for consistency with earlier usage, suggest changing to CWP

  Revised as suggested.

- Figure 9: The caption says that "colors with black dots indicate *significant* differences". Is this correct? I suspect that this should read *insignificant* differences, given that the back dots correspond to when there is near-zero delta-density.

  Thank you for your comment. We confirm that colors with black dots indicate significant differences. It can be observed that most of these black dots are located in darker-colored regions, which represent larger differences in delta-density.

References:

McFiggans, G., Artaxo, P., Baltensperger, U., Coe, H., Facchini, M.C., Feingold, G., Fuzzi, S., Gysel, M., Laaksonen, A., Lohmann, U. and Mentel, T.F.: The effect of physical and chemical aerosol properties on warm cloud droplet activation. *Atmospheric Chemistry and Physics*, *6*(9), pp.2593-2649, 2006.

Heidinger, A. K., Foster, M. J., Walther, A., and Zhao, X.: NOAA Enterprise Cloud Mask Algorithm Theoretical Basis Document (ATBD), NOAA/NESDIS, 2020.