Responses to Reviewers

We appreciate the reviewers for dedicating their time to review our paper and for providing constructive feedback. The page and line numbers we reference to indicate changes in the manuscript correspond to the revised marked-up version.

Reviewer #1

Overview

This study examines how common simplifications in atmospheric models affect predictions of particles that form cloud droplets by comparing a detailed simulation tracking individual particles to simplified approaches. They found that common simplifications can produce seemingly accurate results because different errors cancel each other. Though results seem accurate, seemingly accurate model outcomes mask incorrect physical processes. Incorporating the effect of surfactants on CCN activation is shown to improve predictions, highlighting a pathway to better atmospheric modeling.

The paper addresses an important topic withing atmospheric chemistry: how simplifications of the aerosol loading affect the outcomes of modeled atmospheric processes (namely cloud condensation nuclei activation). This is an ever-present concern as aerosol-cloud interactions remain a large source of uncertainty for future climate change.

The paper makes a concerted effort to explain the impact of oversimplifying particle composition and surface tension in cloud condensation nuclei predictions, demonstrating that while simplified aerosol schemes may produce reasonable CCN predictions, compensating errors are mostly responsible for these results, and ultimately the underlying mechanism is misrepresented.

Their findings that simplified aerosol schemes do produce reasonable CCN predictions, despite oversimplifying a complex aerosol loading, is surprising given what oversimplificiation overlooks. These results provide valuable insight regarding the applicability of simplified schemes, and an updated solution for scenarios when modeling a more complex scheme is necessary. For the most part, the writing is concise and accessible. Some explanations are lacking in details, but overall the paper is well-written and worthy of publication.

Specific Comments #1.1

Figure 1-You don't explain in the caption what the graphics at the top right of each subplot depict. I imagine these are meant to be schematics depicting how individual aerosols are represented in each scenario. Is that correct? If that is correct, shouldn't the effective surface tension in the bottom right plot vary from one particle to the other, similar to subplot (b), which also utilizes an effective surface tension?

We thank the reviewer for pointing out this missing explanation. The graphics in the upper right corner of each panel are schematics of how individual aerosol particles are represented in each scenario. Green denotes organic material and the remaining colors (red, dark-blue, and light-blue) denote the different inorganic species (e.g., sulfate, ammonium, sea salt). We have added an explicit description of these schematics to the Figure 1 caption.

We would like to clarify that these schematics depict particle composition (mixing state), not the surface tension directly. In the particle-resolved cases (a, b), each particle retains both its own size and its own composition, so the organic mass fraction, and hence the effective surface tension, varies strongly from particle to particle. In the composition-averaged cases (c, d), the defining operation is to assign every particle within a given size-range the single bulk-average composition while preserving its original size. This is why the particles in (c) and (d) belonging to the same size range are drawn as identical pie charts of different sizes.

The reviewer raises a valid point that the effective surface tension is not a function of composition alone. In our EST scheme, $\sigma(D)$ is set by the organic surface coverage, which depends on the organic volume

fraction and, more weakly, on particle size through the core-shell (minimum shell thickness) geometry. Therefore $\sigma(D)$ in Comp-EST is not strictly identical across particles. It retains a weak residual size dependence. However, because composition has been homogenized, the particle-to-particle variability is largely removed. PR-EST spans essentially the full range between the organic and inorganic limits (~30–73 mN m^{-1}), whereas Comp-EST collapses onto a narrow band (~65 mN m^{-1}. This near-uniformity is the intended and physically meaningful consequence of composition averaging, and it is precisely the effect that the vertical comparison between panels (b) and (d) is designed to isolate.

Specific Comments #1.2

Line 114 - I think more details are still needed in this section. How are you obtaining and applying the value for fractional surface coverage of inorganic and organic components for each particle? This value will change depending on the particle and amount of surfactant present, so how are you accounting for this variability in this method? Even if these details are provided in Xu et al., 2026 (a previous publication), I recommend including a more thorough description to answer these questions for the reader, especially since this is a key detail in this manuscript.

We thank the reviewer for this suggestion. The fractional surface coverage C_β is not a prescribed constant, it is computed individually for each particle from its own simulated composition. Following Ovadnevaite et al. (2017), we assume the organic material forms a shell (phase β) around an inorganic core (phase α), and that a complete film requires a minimum shell thickness $\delta_{\min}$, so that for a droplet of wet diameter D,

$$C_\beta(D) = \min\left[\frac{V_\beta}{V_\delta(D)}, 1\right], \qquad V_\delta(D) = \frac{4}{3}\pi\left(\left(\tfrac{D}{2}\right)^3 - \left(\tfrac{D}{2} - \delta_{\min}\right)^3\right),$$

where V_β is the particle's organic-phase volume and $V_\delta(D)$ is the volume of a spherical shell of thickness $\delta_{\min}$. Because V_β differs from particle to particle, organic-rich particles reach full coverage ($C_\beta = 1$) while organic-poor particles are only partially covered; and because $V_\delta(D)$ grows as the droplet takes up water, C_β and hence $\sigma(D)$ also vary along each particle's Köhler curve.

We have added Equation (2) and a short description to Section 2.1 of the revised manuscript (P4L109–117) so these details are available without consulting Xu et al. (2026).

Specific Comments #1.3

Line 130 - Where are you getting the data for the mixing state parameter χ from? How are you able to generate these values for an entire region? Where are you getting the data that is allowing you to calculate this? A more thorough explanation of this parameter and how you generate this data should be included.

We thank the reviewer for this question. The mixing state index χ is not an independent input, it is computed diagnostically from the same particle-resolved WRF-PartMC output that underlies all of our analysis. For every computational particle in every grid cell, WRF-PartMC provides the mass of each aerosol species and the statistical weight that the particle represents. For each grid cell we use these per-particle species masses to compute the average per-particle species diversity D_α and the bulk population diversity D_γ, following the definitions of Riemer & West (2013), and then form $\chi = (D_\alpha - 1)/(D_\gamma - 1)$. This yields one value of χ per grid cell, and mapping this over all land grid cells produces the regional distributions shown in the manuscript. In this study the species are grouped into hygroscopic and non-hygroscopic components before computing the diversities.

We have clarified this in Section 2.3 of the revised manuscript (P6L140–142), explicitly stating that χ is derived from the per-particle species masses in the WRF-PartMC output and computed cell by cell to generate the regional maps.

Specific Comments #1.4

Line 130-132 - If χ is only being applied as a diagnostic to interpret prediction errors, how are you obtaining the results such as those in Figure 1? How are you generating a particle-resolved aerosol loading and applying it to predict CCN activation without considering mixing state?

Thanks for bringing up this possible source of confusion. As we mentioned in our reply to previous comments, the mixing state index χ is a summary of how internally or externally mixed the population is. We compute it from the same per-particle data, but it is not an input to the activation calculation. Rather, we use it afterwards to organize and interpret the CCN errors.

Therefore, saying that χ is used "only as a diagnostic" does not mean that mixing state is ignored in the CCN calculation. On the contrary, mixing state is fully resolved in our reference calculation: each particle retains its own dry diameter and full chemical composition, and the critical supersaturation s_c is computed particle by particle from these per-particle properties. The CCN concentration is then obtained by counting the particles whose s_c falls below the environmental supersaturation. No scalar mixing-state parameter or population-averaged composition is used in the particle-resolved activation calculation.

To avoid this ambiguity, we have revised the wording in Section 2.3 (P6L138–141) to state explicitly that χ is a diagnostic computed from the particle-resolved population and does not enter the activation calculation, while the CCN calculations themselves are performed directly on the particle-resolved population with each particle carrying its own size and composition.

Specific Comments #1.5

Line 134 - The term "aerosol data" is rather vague. What specific data are you gathering from the model that are being reported in the manuscript?

We agree that "aerosol data" was too vague. To clarify, all analyses in this study are based on the per-particle output of a single regional-scale, particle-resolved WRF-PartMC simulation. For each simulated particle, this output provides its dry diameter, the mass of each aerosol species, and the number concentration it represents. From these primary quantities we compute all the results reported in the manuscript, including the per-particle hygroscopicity κ and critical supersaturation s_c, the CCN concentrations, and the mixing state index χ. We have revised the opening of Section 2.4 (P6L144–147) to state this explicitly and to remove the vague wording.

Specific Comments #1.6

Line 259-260 - The wording here is very confusing. The language makes it sound like you are only considering a constant surface tension and a particle-resolved aerosol loading in equation 13, but based on the equation itself, you are looking at the effect of incorporating the effective surface tension. I strongly recommend revising this sentence to improve clarity.

We agree the original wording was confusing. As the reviewer notes, Eq. (13) (which became Eq. (14) in the revised paper) is the relative CCN difference between the two particle-resolved cases, PR-CST and PR-EST, so it quantifies the effect of surface tension, i.e., of using CST versus EST, rather than considering CST alone. We have rewritten the sentence in Section 3.4.1 (P11L279–282) to make this explicit:

"Here, $\varepsilon_\sigma^{\mathrm{PR}}$ is the relative CCN difference between the two particle-resolved cases, PR-CST and PR-EST. Since both retain full particle-level composition, it isolates the effect of surface tension alone, i.e., the CCN change caused by using CST (neglecting surfactants) instead of EST (including them)."

Technical Comments #1.7

Line 281 - The introductory phrase should end with a colon, not a period.

We have changed the period to a colon: “Taken together, these results lead to two conclusions:”.

Technical Comments #1.8

Line 354-355 - Because the phrase “composition averaging...surface tension reduction” is an interrupting clause, it should be separated by either em dashes or parentheses, not commas.

We thank the reviewer for catching this. We have rephrased the sentence as below, so that the two simplifications are introduced at the end with a colon, which removes the interruption entirely.

“This study examined how two common simplifications in large-scale aerosol models jointly affect predictions of cloud condensation nuclei (CCN) activation: composition averaging within size ranges and neglect of surfactant-driven surface tension reduction.”

Reviewer #2

Overview

This is a valuable paper. The manuscript addresses an important problem in aerosol-cloud interactions: how common simplifications in large-scale aerosol models affect CCN prediction. The main result is clear and useful. A modal-like treatment using composition averaging and constant surface tension can underpredict CCN relative to a more detailed particle-resolved treatment with effective surface tension, and compensating errors can hide part of that discrepancy.
The paper is strongest when it moves beyond bulk CCN differences to explain the particle-level and size-dependent reasons for the bias. This is a valuable contribution because it shows that reasonable agreement in domain-mean CCN does not mean the activation physics are represented correctly. Overall, this is a solid paper with a practical modeling takeaway.

Comments #2.1

Co-condensation during activation The manuscript may benefit from a brief discussion of organic co-condensation during humidification and activation. If semi-volatile or intermediate-volatility organics repartition as relative humidity increases, particle composition, organic mass, hygroscopicity, and possibly the surface-active material available at the droplet surface may change during activation. This seems relevant to the paper’s central point that composition and surface tension jointly control critical supersaturation.
This does not need to become a new set of simulations, but the authors could acknowledge whether their framework assumes fixed dry-particle composition during activation and discuss how RH-sensitive gas-particle partitioning might affect the magnitude or direction of the reported CCN biases. One relevant recent study is Damha et al. (2023), ”Capturing the Relative-Humidity-Sensitive Gas-Particle Partitioning of Organic Aerosols in a 2D Volatility Basis Set,” *Geophysical Research Letters*, https://doi.org/10.1029/2023GL106095.

We thank the reviewer for raising this point, which is closely connected to our central argument that composition and surface tension jointly control critical supersaturation.

Gas–particle partitioning of semivolatile species is treated dynamically during the WRF-PartMC simulation. However, our CCN calculations begin from a fixed aerosol snapshot. The non-water particle composition at that snapshot reflects the partitioning established under the local temperature, relative humidity,

and gas-phase conditions. During the subsequent activation calculation, this composition is held fixed as the particle takes up water.

Consequently, the additional condensation, evaporation, and repartitioning of semivolatile organic and inorganic species that may occur as relative humidity increases toward activation are not represented. Such processes could alter the amount and composition of solute in the growing droplets, their hygroscopicity, and, where surface-active species are involved, their interfacial properties. Their net effect on critical supersaturation and on the differences among the particle-resolved and composition-averaged treatments is not necessarily unidirectional. It would depend on the identities and volatilities of the semivolatile species, the available gas-phase reservoir, their uptake across particle sizes and compositions, and their partitioning within the droplet.

We have revised Section 2.4 to make the snapshot assumption explicit and to acknowledge this limitation. Quantifying these effects would require coupling droplet activation to dynamic multiphase partitioning and is beyond the scope of the present study.

We have added a corresponding discussion to Section 2.4 (P6L151–157).

Comments #2.2

Clarify CCN versus CDNC interpretation The discussion of the reported error metric could more clearly distinguish CCN concentration from cloud droplet number concentration (CDNC). I suggest adding a simple clarifying sentence when this metric is introduced, such as: ”CCN represents particles that can activate at a specified supersaturation, whereas CDNC represents the number of droplets that actually form in a cloud parcel.” This would help prevent readers from interpreting a percent CCN bias as the same percent bias in cloud droplet number.

We thank the reviewer for this helpful suggestion and have adopted the proposed clarification. When the error metric is introduced in Section 3.4 (P11L253–254), we now add the sentence: “CCN represents the particles that can activate at a specified supersaturation, whereas CDNC represents the number of droplets that actually form in a cloud parcel.” This precedes our existing discussion (based on Ervens et al. (2010)) explaining that the peak supersaturation in a rising parcel adjusts to aerosol loading, so that the reported CCN differences should be regarded as an upper bound on potential CDNC impacts rather than as an equal percentage change in cloud droplet number.

References

Ervens, B., Cubison, M., Andrews, E., Feingold, G., Ogren, J., Jimenez, J., Quinn, P., Bates, T., Wang, J., Zhang, Q., et al. (2010). Ccn predictions using simplified assumptions of organic aerosol composition and mixing state: a synthesis from six different locations. *Atmospheric Chemistry and Physics*, 10(10), 4795–4807.

Ovadnevaite, J., Zuend, A., Laaksonen, A., Sanchez, K. J., Roberts, G., Ceburnis, D., Decesari, S., Rinaldi, M., Hodas, N., Facchini, M. C., et al. (2017). Surface tension prevails over solute effect in organic-influenced cloud droplet activation. *Nature*, 546(7660), 637–641.

Riemer, N. & West, M. (2013). Quantifying aerosol mixing state with entropy and diversity measures. *Atmospheric Chemistry and Physics*, 13(22), 11423–11439.

Xu, X., Curtis, J. H., Yao, Y., Liu, Y., West, M., & Riemer, N. (2026). Role of liquid-liquid phase separation-induced surface tension changes in cloud droplet activation. *Aerosol Science and Technology*, (pp. 1–19).