Response to Reviewer 2
Legend: Reviewer comments - black, author responses - blue, manuscript text changes - green

McMichael et al 2026 ACP:

This article investigates the brightening potential of an MCB like aerosol perturbation to the NEP marine boundary layer. The authors have cleverly combined LES simulations with an offline particle transport model to estimate the Twomey component of the cloud radiative effect. This is an extension of the Idealised Wood 2021 ACP paper and will be of interest to the MCB community. While the authors have explored several subtopics in this idealised study, it would have been better if they had tried to match some of their results with the existing LES studies. The data are all publicly available. This would add more credibility to their model before embarking on an exercise with GCM(s), which I believe is the next logical step. Additionally, it would be useful if the authors could trim certain portions of the results section to enhance the readability of the manuscript. As it stands, it is slightly on the longish end.
My recommendation: publish after the comments below are addressed.

We would like to thank the reviewer for their time, efforts, and helpful comments. We hope that our responses satisfactorily address the reviewer's concerns.

In McMichael et al. (2024), we matched the particle model results to 4 different large-domain LES simulations (204.8 km wide in the cross-plume direction, 25.6 km in the along-plume direction), and demonstrated that the particle model can reproduce LES spreading rates. Outside of this small library of simulations, there are not many simulations using domains large enough to do meaningful particle model-LES comparisons, which was the main motivation for using the particle model to estimate plume spreading rates within our much larger library of background LES conditions.

We agree that it would be beneficial to have a much larger library of large-domain, long duration LES plumes, but each of these simulations can cost hundreds of thousands of core hours (for context, the 4 simulation library from McMichael et al. (2024) cost well over a million core hours), limiting our ability to create a large library at the current time. We believe the existing comparison in McMichael et al. (2024) provides a reasonable validation of the particle model behavior in various environments, but we acknowledge that additional validation simulations would be helpful.

We appreciate the reviewer's suggestion regarding manuscript length. We reviewed the results section and we believe there is one main area where we could shorten the manuscript without substantial loss of information (Section 3.6). We have removed Figures 13 and 14, since the results largely agree with the analysis in Figure 12. We have moved these figures to the Supplement and have edited the discussion in Section 3.6 accordingly. Additionally, the reviewer's comments on Figures 7 and 8 led to new versions of Figures 7, 8, 9, and 10, which

emphasize that the markers on the original plots are ensemble means/medians and the underlying cases do exhibit variability.
Comments:

1. Sec 2.1 and 2.2, can you attribute any physical significance to PC1 and PC2?

Although the PCs are statistical constructs, the Pearson correlation coefficient (R) between each PC and each physical variable provides insight into which variables are most strongly associated with the variability captured by the PCs. The highest R with PC1 is for the change in estimated inversion strength along the trajectory (ΔEIS) and the free-tropospheric mixing ratio (FT q), showing that these two are the most important properties driving variation in PC1. For PC2, EIS and the change in horizontal surface wind speed (ΔWS) are the most important variables, highlighting their influence on the variability represented by PC2. See Fig. 3 in Erfani et al. (2025) and the accompanying discussion for more details.

2. Line 101: quantify "reasonably capture". What does it mean?

At each of the six source locations, nine trajectories corresponding to combinations of -1.5σ, 0, and $+1.5\sigma$ (where σ is the standard deviation) in PC1 and PC2 are selected using the methodology of Erfani et al. (2025). That study then mapped the resulting 54 data points to the corresponding cloud-controlling factors and cloud properties and showed that these 54 trajectories selected from the PC1–PC2 space effectively span the observed ranges of the physical properties and provide a much more representative sampling of the full 1663-trajectory dataset than random sampling.

3. Lines 146 - In LES you are not solving N-S equations. You are solving a filtered form of the N-S equations.

We have added “filtered” to the following sentence:

“As an alternative to explicitly modeling fluid motion and aerosol transport in physical space using the filtered Navier-Stokes equations solved on a numerical grid, as done in the LES,…”

4. Why is the LPM domain size substantially different from that of the LES?

The LPM domain is substantially larger than the LES domain because the two models serve different purposes and have very different computational costs. With the LPM, we can use the LES turbulence statistics and cloud morphology to examine plume evolution over spatial scales that would be prohibitively expensive to simulate directly with LES.

The larger LPM domain allows us to represent changes in plume geometry associated with the evolving mean wind field, including plume elongation and orientation. We state this on line 187, but we modified the following sentence to clarify what we meant by plume geometry changes:

“A 200 x 200 km^2 domain was used to minimize artificial plume self-interaction under periodic boundary conditions on diurnal timescales and to capture changes in plume geometry, such as elongation and changes in plume-axis orientation, associated with ship motion and time-varying mean winds.”

5. Lines 201 to 204: I am not able to understand what the authors are trying to say. Can you rephrase? An equilibrium in what sense?

This section does require some clarification. We are approximating equilibrium “injected aerosol” conditions, assuming a 12-hour active aerosol injection period each day. Different particle decay timescales take different amounts of time to reach peak concentrations that don’t change much between one day and the next (this is illustrated by the figure below). We call this periodic equilibrium, since the domain-mean peak concentration does not vary day-to-day with the periodic injection cycle.

We have edited the following text to hopefully provide clarity:

“To approximate equilibrium aerosol conditions representative of multi-day, repeated injections as expected in an MCB scenario, we repeat the 12-hour daytime aerosol injection until the domain-mean peak aerosol concentration changes by less than 2% between consecutive days. Under this criterion, the number of repeated diurnal cycles necessary for periodic equilibrium for τ_decay = 0.25, 0.5, 1, 2 days is 1, 2, 4, and 8 cycles (days), respectively.”

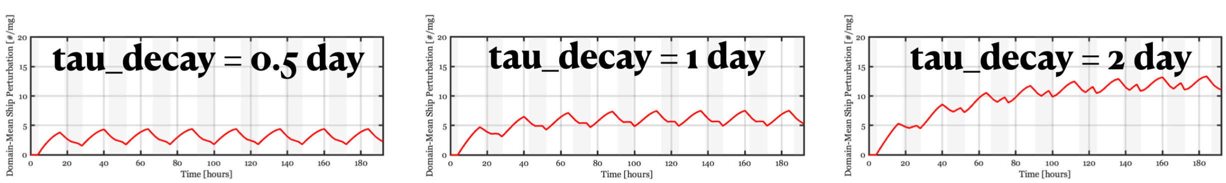

Figure: Domain-mean concentrations as a function of time for repeated daily injections as a function of the particle decay timescale (varying 0.5-2 days).

6. Line 221: Isn't Finc a constant in space? Then why the overbar?

That’s correct. We have removed the overbar and additionally, removed the dependence on x and y, since Finc is only a function of the diurnally varying zenith angle.

7. Lines 229-30: Is it SW radiative forcing or radiative effect? From what I understand it should be a radiative effect. Please clarify and check throughout the manuscript. Forcing is used repeatedly.

Radiative forcing, or Twomey forcing in this study, refers to the radiative flux change attributable to the aerosol-induced cloud albedo change relative to the unperturbed reference state, consistent with the treatment in Wood (2021). We also present the radiative forcing changes as global averages in Figures 7, 9, 10, and 12.

8. Sec 3.1, how are the LWP and cloud fractions defined? Are there any thresholds on LWC or cloud optical depth?

LWP is calculated from the vertically integrated liquid water condensate in each model column. Cloud fraction uses a minimum optical depth threshold of 0.3.

9. Line 317: How can they all influence in nonlinear ways? Can you explain this a bit? It was stated at multiple locations throughout the manuscript.

By "nonlinear" we mean that these turbulence and cloud properties do not influence Twomey forcing independently or in simple proportion to one another. Processes such as turbulence, mean winds, boundary layer depth, and precipitation can alter the spatial distribution of injected aerosol through the spreading rate, while the Twomey response itself depends nonlinearly on local aerosol/cloud droplet concentration. The structure of the overlying cloud field also matters, as the same plume can produce different radiative responses depending on how it overlaps with the cloud morphology, LWP (this would be indirectly accounted for using the LES radiation), and background NAc.

10. Fig 7: is there a way to bring in some uncertainty estimates? Without that I am not sure how to interpret these numbers. Are all of these upper bounds?

The markers represent the ensemble average of all 17 cloudy simulations. We have added IQR shading to the plot, to emphasize that these are ensemble averages. Additionally, we have added IQR shading to Figures 9, 10 and plotted individual cases, medians, and means in Figure 8. We believe the addition of the IQR and case-level data should reduce confusion for readers, as it becomes more clear that the single point is a representative value of the LES/LPM ensemble and not an upper bound or some other ambiguous quantity. While we still want to focus on the regime-specific means (Day 1 versus Day 2) for the discussion, the library variability component is certainly worth demonstrating.

11. Lines 360-370: I am not able to follow how the susceptibility analysis was carried out. Can you please explain it better? Consequently, I am not able to appreciate the importance of Fig. 8.

The confusion may stem from an inaccurate usage of "susceptibility", which is normally reserved for dX/dY-type perturbation studies. This particular analysis is better characterized as a sensitivity analysis, as we are swapping certain Day 2 processes into the Day 1 particle model simulation. We have replaced the "susceptibility" language with "sensitivity" when discussing this section.

We have added two sections of text to hopefully provide clarity and edited the Figure (panel a) to more clearly demonstrate the case-by-case variability.

Section of text in Section 3.3:

"It is not immediately clear which of the meteorological, aerosol, and turbulence variables are most responsible for the weak regime sensitivity. To assess their relative roles, we perform a sensitivity analysis in which Day 2 values of boundary-layer depth, cloud fraction/morphology, particle model turbulence forcing, mean winds, and background aerosol concentration are substituted individually into the corresponding Day 1 simulations for each case, while all other variables retain their Day 1 evolution. We then quantify the resulting change in ΔF relative to the corresponding Day 1 control simulation."

Section of text in conclusions:

"To isolate the effect of turbulence-driven spreading differences on Twomey forcing, we repeated the Day 1 LPM simulations using the Day 2 turbulence characteristics while maintaining the Day 1 cloud morphology, boundary-layer depth, mean wind, and background aerosol concentration conditions. Despite the substantially faster spreading associated with the Day 2 turbulence, the resulting Twomey forcing changes are smaller than 1\%. Therefore, in these simulations, the differences in turbulence-driven horizontal plume spreading alone do not produce appreciable changes in Twomey forcing, even though the underlying aerosol concentration distributions are impacted by the spreading rate differences."

12. How important is the subgrid scale representation of plume tracks in GCMs considering the high uncertainty associated with ACI representation? Would including these subgrid scale effects reduce the uncertainty or increase the uncertainty from an MCB perspective?

Great question. It remains to be seen just how large of an impact the three-dimensional spreading may have on the aerosol size distribution evolution and subsequent activation. However, this study suggests that the horizontal spreading (assuming well-mixed, perfectly activated aerosol) would likely not be of first-order importance in many MCB environments, but getting rid of the current instantaneous spreading rate assumption and instead using a fixed spreading rate would be a relatively simple improvement that could reduce the risk of substantial overcooling.

The purpose of the subgrid-scale plume parameterization would be to reduce structural uncertainty in plume evolution, which could also improve the representation of ACI within the GCM. However, we cannot necessarily conclude that this would reduce the overall uncertainty in simulated MCB forcing, as that would require evaluation against an observational benchmark or a GCM intercomparison study.

13. How would incorporating activation fraction change your results? Have a reservoir of aerosol and activate them when the need arises. This probably is more closer to reality than the current treatment.

We agree that explicitly representing aerosol activation would be more realistic and could potentially be quite important. We are examining this in ongoing work using a three-dimensional particle model to simulate transport to the cloud layer, with this information passed to a parcel model that explicitly represents aerosol activation.

In the case of incorporating a fixed activation fraction, we would expect this to generally scale the magnitude of our current results, without altering the qualitative behavior of the results.

14. Line 453: shouldn't it be divided by TL (TKE/TL)?

TL has units of seconds [s] and TKE has units of [m2/s2], which means the product gives units of [m2/s].

15. The results here seem to suggest that the timing of the aerosol injection is important. Can you explain this a bit? Is it coming from the role played by precipitation? Prabhakaran et al 2023 and 2024 seems to indicate that timing of the aerosol injection doesn't matter for non-precip. clouds. However, Zhang et al 2023 looked at an ensemble of LES studies and concluded that timing matters even in non-precip clouds. Thoughts?

The only pathway for precipitation to impact results in the current study would be through the change in the C_0 parameter in the particle model, which can change the spreading rate. Since we are not coupling the injected aerosols to the cloud microphysics, we cannot comment on the implications of early- versus late-morning precipitation suppression. The primary takeaway from our sensitivity study is that, in the absence of cloud adjustments, the pre-dawn injection allows time for the plumes to spread and reflect sunlight more effectively over the course of the day.