Response to Referee1

We thank Referee 1 for the careful and constructive review and for recognising the value of documenting the tracking implementation in SCALE-SDM. The comments were particularly useful because they identified the same central weaknesses that we also found after re-evaluating the original manuscript: the computational claims were not sufficiently supported, the forward-tracking component was over-framed relative to its algorithmic novelty, and the scientific case study was presented too strongly relative to its purpose as a methodological demonstration.

In response to these comments, and also to related comments from Referee 2, we have made a substantial revision of both the manuscript and the supporting workflow. The revised manuscript differs considerably from the original submission. We changed the title and abstract to emphasise implementation, computational evaluation, and Two-Pass Hybrid Tracking (TPHT). We added controlled computational tests, no-tracking and coalescence-log-only baselines, memory and output-volume diagnostics, SDNC-scaling tests, output-interval sensitivity tests, a sampling-procedure verification, and a three-dimensional TPHT demonstration. The revised abstract now states these changes explicitly (lines 11-29), including controlled computational tests and the 3D TPHT demonstration.

We also changed the framing of the paper. The revised introduction now states that the paper documents a Lagrangian particle-tracking implementation in SCALE-SDM, evaluates wall-clock time, memory footprint, output volume, and scaling relative to no-tracking simulations, and uses the stratocumulus case as a methodological demonstration rather than a claim of new general cloud physics (lines 81-104). We added a distinction between full-lineage backward tracking and target-restricted reconstruction (lines 90-92), and we introduced TPHT as the practical target-restricted strategy for 3D applications (lines 92-95).

Finally, we revised the language throughout the manuscript to reduce overstatement and to avoid claims that the case-study results are new physical mechanisms. We also retained and updated the disclosure that generative AI tools were used only for language editing (line 1034).

Major comments:

1. **Comment:** The paper lacks a sufficient computational evaluation of the algorithms. A thorough evaluation of the memory footprint of both algorithms, the efficiency with which they use computational resources, and their scaling with increasing number of superdroplets would greatly enhance the claims the paper makes along the lines of computational efficiency. Discussing the computational efficiency in the context of the same simulations without tracking would also be extremely beneficial here. Without a more thorough analysis of the computational expense of both algorithms the paper still serves a valid purpose, to comprehensively document the algorithms for the community, but then the expectations and key statements made by the paper should be re-aligned to match this (see specific comments for particular points in the manuscript where this applies, mostly on statements claiming the algorithms are computationally efficient.).

   Reply: We agree. This was the most important weakness of the original manuscript. We therefore added a controlled computational evaluation and substantially reduced unsupported claims of general computational efficiency (lines 538-559).

   The revised manuscript now includes a dedicated computational and storage section (Sect. 3.2.3, lines 537-612). This section reports the machine configuration, compiler and I/O libraries, 3D benchmark settings, wall-clock times, rank-level peak memory, output volume, coalescence-log output, super-droplet number concentration (SDNC) scaling, output-interval sensitivity, and TPHT handoff/storage diagnostics. In particular, the revised manuscript now compares 5% sampled forward and backward tracking cases with a no-tracking baseline and a coalescence-log-only case. The sampled forward and backward tracking configurations showed wall-clock times of 8754-8750 s compared with 8982 s for the no-tracking baseline, and peak memory per rank remained around 649-653 MiB compared with 650 MiB for the baseline (lines 546-550). The main additional burden was diagnostic output volume, not identifier bookkeeping.

We also added Appendix A (lines 875-907), which gives the controlled cold-start 3D benchmark setup and explains how wall-clock cost, output volume, memory use, and instrumented diagnostic timings were measured. This appendix explicitly cautions that values slightly below the no-tracking case should not be interpreted as speed-up, but rather as indicating that sampled-tracking overhead is small relative to run-to-run variability in this controlled setup.

To address scaling, we added Appendix C (lines 973-1004). This appendix reports SDNC-scaling tests with SDNC values of 10, 20, 40, and 80 per cell, using no-tracking, forward-tracking, and backward-tracking configurations. It also reports output-interval sensitivity tests using 30, 60, and 120 s selected-super-droplet output intervals. We explicitly state that these tests describe finite-range empirical behaviour for the tested machine configuration and should not be interpreted as general asymptotic scaling laws.

We now make the limitation explicit. For TPHT, the offline merge/deduplication step was not instrumented as a separate wall-clock metric, so the revised text states that the total end-to-end workflow cost also includes offline parsing, deduplication, repartitioning, and handoff-file writing (lines 577-612). We therefore describe

TPHT as a storage- and diagnostic-workflow demonstration rather than a complete end-to-end performance optimum.

2. **Comment:** It is unclear what the innovations are in the forward-tracking algorithm. I understand the backward-tracking method is particularly fast because of its O(1) lookup of superdroplets. In contrast, to me the forward algorithm looks like a fairly standard approach to selective data output from Lagrangian simulations. Is the method for sampling of superdroplets particularly special? If there is a particular innovation in the forward algorithm it would greatly benefit from being more clearly highlighted. Else, if this algorithm is indeed already a rather standard approach, the paper should make clear what exactly is the goal of documenting it here, e.g. to highlight the particular details of the SCALE-SDM version.

   Reply: We agree with this assessment. In the revised manuscript we no longer present forward tracking as a major standalone algorithmic novelty. Instead, we frame it as an implementation and workflow component that is needed for selected-target output, controlled comparisons, and TPHT forward discovery.

   The revised forward-tracking section now states that the forward algorithm follows selected target super-droplets by assigning persistent identifiers and writing their states at selected-output times (lines 254-255 and 277-281). We then explicitly state that the term "selected target super-droplets" is used rather than assuming statistical representativeness for all diagnostics, and that the forward algorithm should be interpreted primarily as an implementation and workflow component rather than as a claim of general algorithmic novelty (lines 259-261).

   We also revised the sampling description to avoid overclaiming. The shared target-selection and sampling subsystem is now described in Sect. 2.4 (lines 297-346).

The text explains random sampling, stratified sampling, target-count control, bin quotas, fallback behaviour, and the limits of representativeness. In particular, the revised text states that stratified sampling can reduce seed-to-seed variability for diagnostics that depend strongly on height or droplet size, but does not guarantee statistical representativeness for all quantities (lines 335-346).

Finally, we added a practical comparison of backward tracking, forward tracking, and TPHT in Sect. 2.6 (lines 403-464). This section clarifies that the three workflows answer different diagnostic questions rather than being interchangeable. It also directly addresses why simply increasing the number of forward-tracked particles is not equivalent to TPHT: forward tracking remains an a priori selected-output strategy, and increasing the sampling fraction increases output volume with the number of retained particles, output times, and variables written (lines 440-446). TPHT instead separates target discovery from trajectory reconstruction (lines 447-452).

The revised manuscript still includes stratified sampling tests, but it does not claim that this sampling is universally representative. We now state that representativeness must be assessed separately for the diagnostic of interest (lines 434-437 and 911-971).

3. **Comment:** Are the algorithms specific to superdroplets as opposed to any sort of Lagrangian particles/parcels? I suggest rephrasing mentions of superdroplet tracking to Lagrangian particle tracking or simply particle tracking. Even in the title for example “for the Super-Droplet Method” could be removed. If there is something about these algorithms which makes them specific to superdroplets, aside from their implementation in SCALE-SDM, please state this more explicitly. If not, rephrasing superdroplets to Langrangian particles would greatly increases

the scope of this manuscript, widening it's appeal to Lagrangian modellers outside superdroplets and clouds microphysics too.

Reply: We partly agree. We revised the title and terminology to emphasise Lagrangian particle tracking, but we retained "Super-Droplet Method" and "SCALE-SDM" in the title because the implementation, variables, event records, and interpretation are specific to SCALE-SDM and to super-droplets with multiplicity.

The revised title now reads: "A Lagrangian Particle Tracking Framework for the Super-Droplet Method in SCALE-SDM 5.2.6-2.3.1: Implementation, Computational Evaluation and Two-Pass Hybrid Tracking" (lines 1-3). The revised abstract also begins with the more general term "Lagrangian particle histories" while specifying the implementation in massively parallel super-droplet simulations (lines 11-29).

We added text in the discussion clarifying the broader applicability and the model-specific limits. The revised manuscript states that the underlying design is not intrinsically limited to SDM, but requires particle-like entities with persistent identifiers, well-defined ownership during parallel exchange, saved predecessor links, and event records. We also state that the interpretation remains model-specific: in SCALE-SDM, an identifier refers to a super-droplet with multiplicity rather than an individual real droplet (lines 808-814).

We intentionally retained the SDM/SCALE-SDM framing because the paper documents a concrete implementation rather than a fully model-independent software library. We now make this scope explicit rather than presenting the approach as universally ready for all Lagrangian models.

Specific comments:

4. **Comment:** P2 L45-46: consider rephrasing the sentence ending "unique challenges and opportunities". This phrase is vague and not addressed in the paragraph that follows.

   Reply: We removed that wording. The revised introduction now says that applying trajectory/provenance ideas to cloud microphysics requires additional model-specific choices, and then lists concrete implementation issues: identifier definitions, MPI transfer, event logging, output-file organisation, post-processing reconstruction, and computational cost (lines 39-44).

5. **Comment:** P3 L67-69: Rephrase this sentence, at least remove "therefore", but also reconsider the term "computationally efficient" (as per general comment above and because it is somewhat contradictory with the 3-D tracking simulation not being possible due to computational expense).

   Reply: We removed the unsupported general efficiency claim. The revised introduction now says that practical value depends on output structure, memory footprint, I/O cost, and scalability (lines 58-62). The computational evaluation is then described as a main component of the paper (lines 537-612).

6. **Comment:** P3 L80-83: Reconsider description of the tracking algorithms here as "robust", "computationally efficient" and "designed for broad scientific inquiry"; it is not clear how the algorithms presented are robust or broad, please either explain what is meant by these terms somewhere in the manuscript or rephrase this. The term computationally efficient is misleading (as above), and this is also the case in P19 L601.

Reply: We rewrote this framing. The revised text now describes a documented, scalability-aware implementation and then explicitly states the components of the framework and the evaluation metrics (lines 75-90). Summary and conclusions now use more limited wording: the study documents and evaluates a Lagrangian particle-tracking framework and presents TPHT as target-restricted reconstruction, not a universal solution (lines 831-836).

7. **Comment:** P3 L88: I have trouble with the phrase “statistically robust” to describe the forward tracking algorithm. On P12 L360-361 a 5% sampling is used and it is not immediately obvious, nor shown statistically that this corresponds to a robust sampling of the population. On P19 L572 the term “statistically representative” is likewise problematic.

   Reply: We removed these claims. The revised manuscript uses "selected target super-droplets" rather than "representative cohort" in the forward-tracking section (lines 254-262). Sect. 2.4 now states that representativeness depends on sampling fraction, binning variables, target diagnostic, cloud regime, simulation duration, and output interval (lines 335-340). Sect. 4.2 and Appendix B also state that the sampling tests are verification tests, not proof of 3D representativeness (lines 787-791 and 909-971).

8. **Comment:** P3 L91-93: “Other researchers . . . ” please be much more specific, consider using one or two example(s) which may be more effective at conveying what is meant than such a general sentence. Similar point for the sentence P18 L592-595.

   Reply: We added specific examples. Matsushima et al. (2023) is now discussed as an example of simple forward identifiers used in high-resolution SDM simulations (lines 63-74). We added Telford (1955), Dziekan and Pawlowska (2017), Hoffmann et al. (2017), Li et al. (2022), and Lim and Hoffmann (2023, 2024) to clarify prior work on lucky droplets, stochastic growth, and event/history-based Lagrangian diagnostics (lines 75-80 and 414-420).

9. **Comment:** P4 L98: I find "presenting key scientific insights" misleading or overstated, consider rephrasing or being more specific.

   Reply: We changed the framing. Section 3 now begins by stating that the purpose is to demonstrate the TPHT workflow, not to establish a new warm-rain mechanism (lines 466-467). The summary also states that the case study is a methodological demonstration rather than a new cloud-physics interpretation (lines 857-863).

10. **Comment:** From what I understand the backward algorithm creates a very large number of separate small files. Is this correct? And what is the structure of these files and was there any consideration on how big/small they should be to optimise reading and writing them? Perhaps memory can be saved by reconsidering how these files are written/organised too.

    Reply: We added and reorganised the output-description text. Figure 1 and its caption (lines 213-230) now explain the separation between SD_selected_NetCDF_* selected-super-droplet files and SD_coal_output_NetCDF_* coalescence-event files (lines 185-192). After Algorithm 3, the revised manuscript summarises the roles of SD_selected_NetCDF_, SD_coal_output_NetCDF_, raw .ids streams, and deduplicated .ids handoff files (lines 396-402). The TPHT configuration table clarifies which pass writes or reads each output type (lines 531-536). Storage numbers are now discussed in Sect. 3.2.3 (lines 584-592).

11. **Comment:** If the pre_sdid and pre_dmid are also written to file, please add them to Fig. 1.

    Reply: Done. Figure 1 caption (lines 213-230) now states that forward tracking writes sd_id and dm_id, while backward tracking writes pre_sdid and pre_dmid pointing to the predecessor record at the previous saved output time (lines 170-173).

12. **Comment:** It may be worth already mentioning in the backward algorithm's description that it is expected to have a very large memory footprint given how much data it outputs, as is then proven later in the manuscript for the DYCOMS II test-case.

    Reply: We now state this at multiple points. The backward algorithm section says that output frequency and target-set size still control storage and I/O (lines 205-211). Sect. 2.6 discusses full-population backward tracking and storage cost (lines 421-426). Sect. 4.2 makes storage and I/O the first limitation (lines 764-780).

13. **Comment:** In the forward algorithm (P9 L260), what are the array dimensions? If they are something like (Time, Domain ID, Superdroplet ID) they would be very memory intensive and could be better optimised, perhaps this too could be pointed out.

    Reply: The revised forward post-processing section was simplified. It now describes direct aggregation over output times and MPI-rank files, without implying a dense Time x Domain x Super-droplet array (lines 290-296).

14. **Comment:** On P18 L570 you might consider data compression, potentially with

AI, as a useful strategy to pursue in the future in order to reduce the algorithms' memory costs.

Reply: We added compression and automatic cleanup as future workflow optimisations for TPHT raw ID streams (lines 800-802). The conclusion also mentions compression or automatic cleanup of temporary raw-ID streams as future work (lines 867-871).

15. **Comment:** The lookup of superdroplets is O(1) and so I believe this means it is no longer the most critical scaling of the algorithm. Does the complete algorithm hence scale linearly with superdroplets? Or perhaps in some way with the batch size of the post-processing jobs?

    Reply: We clarified that $O(1)$ applies to a single predecessor query, not to the complete workflow. The backward algorithm section states that total cost still depends on selected targets, output times, variables written, and files read (lines 139-142). The post-processing section repeats this point (lines 240-245). Sect. 3.2.3 states the same for TPHT reconstruction (lines 605-607). Appendix C gives empirical SDNC and output-interval sensitivity tests (lines 973-1003).

16. **Comment:** On P7 L212-213: the statement after " to further optimize" is not an optimisation, it is a description of ordinary cache-loading in computers. Please reconsider this statement.

    Reply: We revised the wording. The text now says that trajectories can be processed in batches and recently accessed files can be cached temporarily to limit memory use and repeated disk access (lines 240-245).

17. **Comment:** P13 L401: Can you estimate the memory footprint of the coalescence logs that would've been produced if coalescence had been activated? It seems important information for a user who would use the algorithm with coalescence enabled. Or perhaps this could be clarified by a more detailed discussion of the files the algorithms output as commented on above.

    Reply: We added a coalescence-log-only benchmark. Sect. 3.2.3 reports that the coalescence-log-only case wrote about 4.37 MiB of coalescence-event diagnostics (lines 549-550). Appendix A decomposes output volume and notes that the sampled tracking cases produced about 121 MiB of scientific output, of which about 5% was coalescence-event log output (lines 891-899). The TPHT case also reports linked coalescence records in Sect. 3.3 (lines 635-639).

18. **Comment:** a call to ADD_RANDOM_PERTURBATION_ALL occurs which doesn't occur in the backward algorithm. Why is this? Consider stating it is not essential to the algorithm if it is not, or explain why is it necessary for the forward (but not backward) algorithm.

    Reply: This implementation-specific and nonessential step was removed from the revised high-level forward algorithm. Algorithm 2 now includes standard SDM physics/motion, identifier propagation, coalescence-event logging, TPHT ID output, and selected-SD output, but not ADD_RANDOM_PERTURBATION_ALL (lines 283-289).

19. **Comment:** Sorting is mentioned in post-processing, however non-specific sorting algorithms are often extremely costly and so the choice and cost of the algorithm should be elaborated on if the computational performance of the algorithm is further presented (as commented on above).

Reply: The forward post-processing description was revised. It now describes direct aggregation of selected-output records and no predecessor-link lookup (lines 290-296). Computational cost is addressed through the controlled benchmark and Appendix C rather than by emphasising a sorting-heavy Python workflow (lines 544-550 and 974-1003).

20. **Comment:** Please elaborate on the author contributions, CL is not even mentioned

Reply: Done. The revised Author Contributions section now states: "CY, SS, and CL designed the study and the numerical experiments. SS and CY developed the model code and CY carried out the simulations, post-processing, and analysis. CL provided scientific guidance on the experimental design and interpretation of the results. CY prepared the manuscript with contributions, comments, and revisions from SS and CL." (lines 1028-1034)

Technical comments:

Grouped response for corrected, deleted, or no-longer-applicable technical points

We thank the reviewer for the detailed technical corrections. Many of these points became no longer applicable because the revised manuscript was substantially restructured: the original forward/backward case-study comparison, the original physical-result figures, and the quasi-two-dimensional full-lineage backward experiment were removed or replaced. The revised manuscript now uses a new title, a new abstract, revised algorithms, a TPHT workflow, controlled benchmark appendices, and different main figures.

The following technical points were corrected or removed in the revised manuscript.

First, vague or overstrong wording was removed or moderated, including “Large-Eddy Simulation model”, “compelling scientific outcomes”, “robust”, “statistically robust”, “novel insights”, “significant pathway”, and “impossible to uncover”. The revised manuscript describes the case study as a methodological demonstration and avoids presenting the results as new general cloud-physics mechanisms.

Second, the Introduction was revised to avoid unnecessary quotation and to cite Matsushima et al. (2023) more directly as an example of particle identifiers used in SDM diagnostics, including reference to their Fig. 6. The phrase “The importance of this contribution is significant” was removed (lines 66-67).

Third, acronym and model-version wording was revised. SCALE-SDM version 5.2.6-2.3.1 is now stated in the model description, and the use of “the SDM” was checked and adjusted where appropriate.

Fourth, formatting and terminology issues in the algorithm sections were corrected. The unnecessary paragraph break was removed, “antecedent” was replaced with clearer wording, repeated text was removed, and “memory usage and I/O bottlenecks” was rephrased to avoid implying an artificial balance between the two. The typo in which “pre_dmid” appeared twice was corrected to “pre_dmid” and “pre_sdid”.

Fifth, the old database and sorting discussion in the forward post-processing section was removed. The revised forward post-processing description now states only that analysis scripts aggregate selected-output records over output times and MPI-rank files. No sorting-heavy workflow is presented as part of the algorithmic contribution.

Sixth, the old eddy-hopping definition and related wording were removed. Eddy hopping is now only mentioned as an example of a previously studied diagnostic target class, with citations to Cooper (1989), Grabowski and Abade (2017), and Abade et al. (2018), rather than being defined in a way that could misrepresent the process.

Seventh, old experiment-table and figure inconsistencies were removed by restructuring the experiments. The previous Table 1 issue involving the coalescence setting is no longer applicable. The revised Table 1 is a workflow-comparison table, and Table 2 now describes the TPHT forward-discovery and backward-reconstruction passes.

Eighth, old figure-specific issues involving the original Figs. 2, 4, 5, 7, and 8 are no longer applicable because these figures were removed or replaced by new benchmark, TPHT handoff/storage, target-diagnostic, and TPHT trajectory figures. The new figures use minutes consistently where appropriate.

Finally, the manuscript language was revised throughout to reduce overly bold or AI-like phrasing. The acknowledgement now discloses the use of Gemini and ChatGPT for language editing only.

Technical comments that merit separate replies

1. **Comment:** P8 L 230: does the INVALID ID consume memory? If so, is it transferred by MPI communication? Both these things could be potential sources of (minor) optimisation of the forward algorithm.

   Reply: We clarified the role of INVALID_ID in the revised algorithm descriptions. Non-target or invalid super-droplet slots are assigned INVALID_ID (lines 154-155 and 263-266), and selected target identifiers are exchanged with super-droplet state variables during MPI boundary transfer (lines 259-260). INVALID_ID is retained as a sentinel value in the identifier arrays so that target and non-target records can be distinguished consistently during output and MPI exchange. This introduces a small per-record identifier-field overhead, but it avoids ambiguity in distributed output and post-processing. In the sampled benchmark, this overhead was small compared with output writing; identifier bookkeeping and boundary-tracking exchange were much smaller than selected-super-droplet output writing (lines 891-899).

2. **Comment:** P11 L353: please explain how momentum exchange between super-droplets and the fluid is taken into account, or provide a reference.

   Reply: We added references to the SCALE-SDM formulation. The revised text now states that momentum exchange between super-droplets and the fluid is taken into account following the formulation described by Shima et al. (2009, 2020) (lines 503-504).

3. **Comment:** P13 L398: Why are there more output intervals for the backwards than forwards tracking?

   Reply: This issue is no longer applicable because the original unequal forward/backward experiment comparison was replaced. The revised controlled benchmark uses consistent sampled forward and backward tracking configurations with the same selected-super-droplet output interval (Appendix A, lines 875-907). The TPHT experiment intentionally uses different outputs for the forward-discovery and backward-reconstruction passes because they have different roles: the forward-discovery pass writes raw .ids streams and suppresses selected-SD output, while the backward-reconstruction pass writes selected target histories every 10 s (lines 523-525 and Table 2, lines 531-536).

4. **Comment:** P5 L144: please define a "sequence number". In the same paragraph it would also be helpful to the reader to explicitly state somewhere that "pre_sdid" stands for pre superdroplet identity and "pre_dmid" stands for pre domain identity.

   Reply: We revised the identifier description. The manuscript now avoids the ambiguous term "sequence number" and instead defines the predecessor-link pair in terms of the saved output record. In the output files, pre_dmid and pre_sdid together identify the predecessor record at the previous saved output level: pre_dmid denotes the MPI rank containing that predecessor record, and pre_sdid denotes the local record index within the corresponding rank-local output file (lines 156-158). Figure 1 and its caption also distinguish the identifiers written by forward tracking and backward tracking (lines 213-230).

5. **Comment:** P5 L145: P5 L154: Is the if_coal flag always set to 1? Also in Fig. 1 if_coal = 1 makes it seem like it is always 1?

   Reply: We clarified this point in both the text and Figure 1 caption (lines 190-192 and 213-230). The if_coal variable is a binary interval marker. It is initialised to 0 for selected target records and reset to 0 after each selected-super-droplet output. It is set to 1 only if at least one recorded coalescence event involving that selected target occurred during the current output interval. Figure 1 shows an example in which such an event occurred; it is not intended to imply that if_coal is always 1.

Final note on minor technical and typographical corrections

Several additional comments referred to sentences, figures, or experiments that no longer exist in the revised manuscript. These include the old forward-tracking result section, the old quasi-two-dimensional backward case, the previous physical-diagnostics figures, and the original table of experiments. Where the affected material was retained, we corrected the wording, citation formatting, terminology, figure references, and captions as requested. Where the affected material was removed, the comment is no longer applicable to the revised manuscript. We nevertheless checked the corresponding issue throughout the revised text to avoid carrying over the same problem.

Reference

Abade, G. C., Grabowski, W. W., and Pawlowska, H.: Broadening of cloud droplet spectra through eddy hopping: Turbulent entraining parcel simulations, Journal of the Atmospheric Sciences, 75, 3365–3379, https://doi.org/10.1175/JAS-D-18-0078.1, 2018.

Cooper, W. A.: Effects of Variable Droplet Growth Histories on Droplet Size Distributions. Part I: Theory, Journal of Atmospheric Sciences, 46, 1301–1311, 1989.

Dziekan, P. and Pawlowska, H.: Stochastic coalescence in Lagrangian cloud microphysics, Atmospheric Chemistry and Physics, 17, 13509–13520, https://doi.org/10.5194/acp-17-13509-2017, 2017.

Grabowski, W. W. and Abade, G. C.: Broadening of cloud droplet spectra through eddy hopping: Turbulent adiabatic parcel simulations, Journal of the Atmospheric Sciences, 74, 1485–1493, https://doi.org/10.1175/JAS-D-17-0043.1, 2017.

Hoffmann, F., Noh, Y., and Raasch, S.: The Route to Raindrop Formation in a Shallow Cumulus Cloud Simulated by a Lagrangian Cloud Model, Journal of the Atmospheric Sciences, 74, 2125–2142, https://doi.org/10.1175/JAS-D-16-0220.1, 2017.

Li, X.-Y., Mehlig, B., Svensson, G., Brandenburg, A., and Haugen, N. E. L.: Collision Fluctuations of Lucky Droplets with Superdroplets, Journal of the Atmospheric Sciences, 79, 1821–1835, https://doi.org/10.1175/JAS-D-20-0371.1, 2022.

Lim, J.-S. and Hoffmann, F.: Between broadening and narrowing: How mixing affects the width of the droplet size distribution, Journal of Geophysical Research: Atmospheres, 128, e2022JD037900, 2023.

Lim, J.-S. and Hoffmann, F.: Life cycle evolution of mixing in shallow cumulus clouds, Journal of Geophysical Research: Atmospheres, 129, e2023JD040393, https://doi.org/10.1029/2023JD040393, 2024.

Matsushima, T., Nishizawa, S., and Shima, S.: Optimization and sophistication of the super-droplet method for ultrahigh resolution cloud simulations, Geosci. Model Dev. Discuss., 2023, 1–53, https://doi.org/10.5194/gmd-2023-26, 2023.

Shima, S., Kusano, K., Kawano, A., Sugiyama, T., and Kawahara, S.: The super-droplet method for the numerical simulation of clouds and precipitation: A particle-based and probabilistic microphysics model coupled with a non-hydrostatic model, Quarterly Journal of the Royal Meteorological Society, 135, 1307–1320, https://doi.org/10.1002/qj.441, 2009.

Shima, S., Sato, Y., Hashimoto, A., and Misumi, R.: Predicting the morphology of ice particles in deep convection using the super-droplet method: development and evaluation of SCALE-SDM 0.2.5-2.2.0, -2.2.1, and -2.2.2, Geosci. Model Dev., 13, 4107–4157, https://doi.org/10.5194/gmd-13-4107-2020, 2020.

Telford, J. W.: A new aspect of coalescence theory, Journal of Atmospheric Sciences, 12, 436–444, 1955.