

RC1:

We are very grateful for your insightful and constructive comments, which have significantly improved the manuscript. We have revised the paper thoroughly, including careful language editing and a number of substantive changes. Below, we provide our responses to your specific comments.

Major comments

The development of an emulator of the K15 model is a potentially significant advance for the community. However, quantification of the accuracy of this emulator is currently lacking. The authors state on lines 229 to 234 that the fit looks “very good” and that, even in the worst cases, “the fit is decent”; these are however fundamentally value judgements. I recommend that the authors pick a representative set of, say, 100 cases (including some extremes) and quantify the degree to which disagreement between the emulator and the parent model propagates to disagreement in their simulated impacts (i.e. contrail lifetime and EF, not just inputs to CoCiP).

Response: We thank the reviewer for this helpful suggestion and agree that the original manuscript did not sufficiently quantify how disagreement between the emulator and the modified K15 model propagates to the final simulated contrail impacts. To address this, we performed additional simulations in which the modified K15 model was implemented directly in CoCiP (pycontrails) for a representative subset of lubrication-oil size distributions, including cases spanning the range considered in the manuscript. The resulting comparison, including the effect on the simulated contrail metrics, is now presented in Supplement S6, and is referenced in lines 265-269 of the revised manuscript.

On a similar basis, it would be useful to quantify the degree of (dis)agreement between results from CoCiP with the standard and K15 models when simulating fossil jet fuel contrails (lines 256-258). While S6 does show how the two models differ in terms of activation fraction, this difference is not propagated to an overall impact; furthermore the different formatting choices and lack of parity plotting between the results in S6 make it very difficult to establish whether there really are meaningful differences between the two.

Response: We thank the reviewer for this useful comment. We agree that Supplement S6 did not clearly convey the significance of the differences between CoCiP’s standard soot activation function and the modified K15 model for fossil jet fuel contrails. To improve both the quantification and the visualization of these differences, we implemented the modified K15 soot activation directly in CoCiP using soot properties corresponding to those assumed in K15 and propagated the differences to the simulated contrail cirrus outcomes. We now include these results in Supplement S7 as a box-and-whisker plot of the percentage difference across all simulated contrail cirrus cases. We also refer to these results in line 302-303 of the revised manuscript.

I was surprised that homogeneous nucleation of water was not discussed. Given the potential for extremely high RH in the early plume of a hydrogen-fueled contrail, Zink et al. argue in their preprint (Contrail formation for aircraft with hydrogen combustion - Part 1: A

systematic microphysical investigation) that homogeneous droplet nucleation (HDN) may be significant under some circumstances. It is not clear to me whether neglecting HDN is or is not likely to significantly change the conclusions of this study, but I would recommend that the authors at least clearly highlight whether or not this eventuality is covered by their modelling approach (since the base model by Kärcher et al. was developed for kerosene-based contrails and therefore does not appear to include HDN). Whether or not it is, I would also recommend at least a brief discussion of the implications of HDN for their study since this mechanism would be expected to be possible even in the absence of lubrication oil.

Response: Zink et al. found that HDN could affect ice particle numbers when aerosol concentrations are low or absent. However, they also argued that, under such conditions, HDN would likely not produce higher ice particle numbers than in cases with medium to high ambient aerosol concentrations (Zink et al., 2026). Based on the results of Zink et al. we therefore argue that HDN would be unlikely to alter our results, since we consistently assume high ambient aerosol concentrations throughout our analysis. For this reason, we do not include HDN explicitly, nor do we extend the K15 framework to represent it, as HDN is not covered by the K15 model. To make this reasoning clearer, we have added a section to the main manuscript (lines 181-188) explaining the basis for its omission.

The authors make recommendations regarding engineering decisions which, while plausible, are not fully justified by this work. A recommendation that future engine design should avoid venting of lubrication oil into the main exhaust (line 472) seems premature given that only one aircraft/engine combination is assessed and given that there is as-yet almost no data on how difficult it might be to re-engineer an engine to avoid this (and therefore whether there may be efficiency penalties). I would therefore suggest that statements such as that in closing – that lubrication oil handling “should” be treated “as a critical parameter” – be instead tempered. A defensible statement might be that, for example, a 10 times reduction in lubrication oil emissions in the exhaust appears to reduce hydrogen contrail RF by 60% compared to a representative base case (this latter number read from figure 5 – clearly a better assessment is needed). Whether or not this is critical depends on balancing that against the alternatives, including the potential costs of doing so (which are not evaluated here).

Response: We thank the reviewer for this important comment and agree that our original statements regarding engineering implications were too strong given the scope of the present study. While our sweep over lubrication-oil size distributions was intended to partly capture the effects of different leakage and/or venting conditions, potentially associated with different engine designs and temperatures, we do not explicitly assess the feasibility, trade-offs, or potential efficiency penalties associated with such design changes. We have therefore tempered the corresponding statements in the revised manuscript and now restrict the discussion to what is directly supported by our results (lines 586-587).

Minor comments

The authors introduce a well-reasoned set of caveats in the discussion, which I found very encouraging. However, these caveats did not propagate to the abstract which surprised me. Given that the authors appear to have identified potentially significant limitations, I would

recommend that the abstract include at least a statement to the effect that these findings need to be validated by both a more comprehensive modelling campaign (including e.g. a higher-fidelity aircraft representation, alternative meteorological/contrail models, and a larger dataset of flights) and of course measurements given that there are as-yet no fully hydrogen-fueled aircraft in service or testing.

Response: We appreciate this suggestion and have revised the manuscript accordingly. The caveats identified in the discussion should also be reflected in the abstract. We have therefore added sentences to the abstract noting that the present findings require further validation through more comprehensive modelling and, ultimately, through measurements once hydrogen-fuelled aircraft data become available (lines 14-17).

It is often unclear whether the authors are referring to CoCiP (essentially a modelling approach, given that the original Fortran codebase is not – to my knowledge – publicly available) or pycontrails (a public, Python-based implementation of CoCiP). If the latter, as I expect given line 274, I would recommend that the authors specifically refer to “CoCiP as implemented in pycontrails”. For example, line 264 states that “CoCiP has a built in aircraft performance model”. From what I understand however, pycontrails includes an implementation of the Poll-Schumann model which it uses to initialize its implementation of CoCiP (but which can equally be used to initialize other models, such as a dry advection code).

Response: We thank the reviewer for this helpful comment and agree that the distinction between CoCiP and its implementation in pycontrails should be made clearer. We now state explicitly near the beginning of the manuscript that we use CoCiP as implemented in pycontrails (line 88), and that subsequent references to CoCiP refer to this implementation unless otherwise stated (lines 148-150). We have also adjusted the wording in the relevant passages to reduce ambiguity.

Similarly, I was surprised to see that the gridded version of CoCiP was used (line 274) but that the paper describing that code (Engberg et al., 2025) was not cited.

Response: We thank the reviewer for identifying this omission. We now cite Engberg et al. (2025), which describes the gridded version of CoCiP, at the point where this model configuration is first introduced in the manuscript (lines 148–150).

A small request: the zonal plots (e.g. Figures 2, 3, 4) are rather confusingly oriented. It is more conventional for altitude to be the vertical axis, and latitude the horizontal. To ease reader interpretation I would recommend changing these figures accordingly.

Response: We thank the reviewer for this helpful suggestion. We have revised all zonal plots so that altitude is now shown on the vertical axis and latitude on the horizontal axis throughout.

There are some minor typographical, grammatical, and formatting errors throughout the document (e.g. malformed citation on line 71; “is” instead of “are” on line 196; “effects” instead of “affects” on line 195; missing space on line 74; incorrectly formatted citations on

lines 86, 136, 138, 155, 207...). I would recommend that the authors review the document thoroughly to eliminate such errors.

Response: We thank the reviewer for pointing out these typographical, grammatical, and formatting issues. We have carefully proofread the revised manuscript and corrected the errors identified by the reviewer, along with additional minor inconsistencies found during the revision.

Finally, I would suggest that the authors aim to make their assessment more quantitative. For example, section 4.2.1 compares different lubrication oil parameters but provides almost no quantitative assessment; almost all statements refer to high, low, small, large. It would be of great benefit to the reader to understand, for example, what the average EF per flight meter was across the 20 flights in the baseline case, and the percentage reduction (or increase) which was achieved when testing different factors. Such quantitative assessments enable subsequent researchers to compare their own models directly, and therefore provide outsized value in terms of advancing the field. If the authors could therefore augment their analysis with more quantitative evaluation (not just in 4.2.1 but throughout) I believe it would improve the utility of the manuscript for the community.

Response: We thank the reviewer for this valuable suggestion and agree that a more quantitative presentation improves the utility of the manuscript. We have therefore revised the text throughout to include quantitative estimates of the simulated changes, rather than relying mainly on qualitative descriptors. For example, in Sect. 4.1 we now report that the base case yields an average reduction in EF_{pfm} of 66% relative to fossil jet fuel (lines 388–389), and we have added the corresponding base-case and fossil-jet-fuel values to Fig. 5 to support direct comparison. Similar changes have been made throughout the manuscript where appropriate.

RC2:

We sincerely thank you for your thoughtful and constructive comments, which have helped us to improve the manuscript substantially. In response, we have carefully revised the text throughout and made a number of substantial changes to the manuscript. Below, we provide our responses to each of your comments.

Major comments:

1. Since the role of lubrication oil particles is uncertain, the authors explore many possible scenarios. Their assumptions about oil emissions and particle sizes are based on measurements behind conventional aircraft and laboratory experiments, since no such measurements exist yet for hydrogen engines.

In my view, however, not all potential scenarios are fully covered.

a) The smallest geometric mean radius of oil particles investigated is 5 nm. However, volatile

particle sizes can reach down to ~ 1 nm (e.g., Figure 3 in Yu et al. 2024). At present, it is unclear whether such small particles could form from evaporated oil at cruise altitude (and measurement devices may struggle to detect them). Still, this is one potential scenario missing in the analysis. For the same oil mass emission index, the particle number increases by more than a factor of ~ 100 when the size changes from 5 nm to 1 nm. Of course, smaller particles are more affected by the Kelvin effect, so not all may activate. But given the high supersaturations in hydrogen combustion, I can still imagine that the larger number of oil particles could produce ice crystal numbers exceeding those from conventional kerosene combustion. For completeness, such “worst-case” scenarios should be included.

Response: We thank the reviewer for this important suggestion. We have now extended the analysis to include lubrication-oil droplet size distributions with geometric mean radii of 1 and 3 nm. The corresponding results have been added to all panels of Fig. 5 (and Fig. 6). This addition broadens the scope of the study and provides further insight into the sensitivity of the results to very small particle sizes.

Including these cases also helped to clarify limitations of the K15 and modified K15 models. In particular, for very small mean radii, quenching can occur before supersaturation becomes sufficiently high for substantial activation, a behaviour that is especially pronounced in these cases. This likely leads to an underestimation, rather than an overestimation, of the number of ice particles formed. We have therefore added discussion of this limitation of the K15 model for small particle sizes in Sect. 3.1.1 (lines 173–179), Sect. 4.2 (lines 436–453), and the Discussion (lines 545–552).

b) On the other hand, the authors should also emphasize the “best-case” scenario more strongly. As discussed, the impact of oil could be significantly reduced—or even eliminated—if venting does not occur in the hot sections. In this ideal case, ice crystals would only form on ambient aerosols, which cannot be mitigated through technical measures. When estimating the radiative effect of contrail cirrus, the natural variability of aerosol properties should be represented.

Currently, the authors assume a constant ambient aerosol number concentration of 1000 cm^{-3} . They argue that Bier et al. (2024) already investigated the effect of ambient aerosols, which is why the current study focuses on lubrication oil properties. However, Bier et al. (2024) only addressed the formation stage and did not estimate the radiative effect of contrail-cirrus.

In the current study, ERA5 reanalysis data for the year 2019 are used to represent background conditions for temperature, pressure, and relative humidity. Ideally, an external aerosol module should also be included to realistically capture the large variability of aerosol properties. At the very least, the assumed aerosol number concentration should be varied within realistic ranges to better assess how sensitive the radiative effect is to ambient aerosol variability.

Response: We appreciate this suggestion and have revised the manuscript accordingly. The reviewer is correct that Bier et al. (2024) considered only the formation stage and did not assess radiative effects, so our original justification was incomplete. We have therefore added a sensitivity analysis of ambient aerosol number concentration in Sect. S8 of the Supplement, and now discuss these results, together with the corresponding “best-case” scenario in Sect. 4.3 (lines 497–509), and in the Discussion (lines 576–582).

We agree that ambient aerosol variability should ideally be represented more realistically when estimating contrail cirrus radiative effects. However, a full treatment of this variability would require a broader modelling framework than adopted here and lies beyond the scope of the present study. We therefore highlight it as an important area for future research, while noting that the present manuscript focuses primarily on lubrication oil and on conditions in which lubrication-oil-derived droplets dominate contrail particle formation.

2. Figure 1: I do not understand the physical basis of the activation fraction dependency on ambient temperature. The decrease in activation fractions of lubrication oil particles with increasing temperature seems plausible to me. However, the increase in activation fractions of ambient aerosols with temperature is less intuitive. In lines 347–350, you argue that with lower activation of lubrication oil particles, it takes longer to quench the plume supersaturation, which allows more ambient aerosols to be entrained and activated at higher temperatures. If this competition effect were the explanation, I would expect a clear dependence on the number of oil particles. Yet, the activation fractions of ambient aerosols appear very similar across all panels.

Could you check whether this temperature dependence still exists when the number of oil particles is set to zero? If it does, this would seem to contradict the findings of Bier et al. (2024), who showed that the number of ice crystals formed decreases with increasing temperature due to shorter-lasting and lower supersaturations.

Response: We thank the reviewer for this important comment. Our original discussion did not clearly distinguish the different mechanisms affecting quench time. While larger numbers of lubrication-oil droplets do reduce quench time and aerosol entrainment, this effect is smaller than the temperature dependence associated with faster supersaturation depletion at lower temperatures. Lower temperatures lead to higher supersaturations, which in turn increase the condensational sink in K15 and shortens the quench time. This aspect of the model, which was previously not described clearly enough, is now discussed in Sect. 3.1.1 (lines 173–179), Sect. 4.2 (lines 436–453), and the Discussion (lines 545–552).

We have also confirmed that this temperature dependence remains present when lubrication-oil emissions are set to zero. This shows that the behaviour is a consequence of the quench treatment itself, rather than being driven primarily by competition with lubrication-oil particles. A similar feature can also be seen for ambient aerosols in combination with soot in earlier work with K15, for example in Fig. 2 panel (c) of (Bier & Burkhardt, 2019).

3. It could be helpful to compare the ice crystal numbers obtained here with those reported by Bier et al. (2024) and to discuss possible reasons for any differences. In addition, Zink et al. (2025) also studied the competition between lubrication oil particles and ambient aerosols in hydrogen combustion. A comparison with their results might provide further insight.

We thank the reviewer for this helpful suggestion. We now compare and discuss our results in relation to (Zink et al., 2025) in the Discussion (lines 557–563). A fully quantitative comparison was not possible, however, because of the substantial differences in modelling approach. We also compare the reduction in ice crystal number obtained for our base case with the results of (Bier et al., 2024) in the Discussion in lines 514-516.

4. The authors argue that the decisive factor for ice crystal formation is the number of particles present. However, this conclusion cannot be directly inferred from Figure 5. When the results are plotted as a function of geometric mean radius, it is not immediately clear how radius influences the oil particle number (for a fixed oil mass emission index) versus how radius affects activation suppression due to the Kelvin effect. I therefore strongly recommend plotting the results as a function of oil particle number N_{oil} (per flight meter). In this case, the highest N_{oil} would correspond to the smallest radius, and the lowest N_{oil} to the largest radius. In the second row, the two impacts of oil particle radius would then be disentangled, making it possible to directly see how many of the oil particles activate into ice crystals.

Response: We thank the reviewer for this excellent suggestion. Presenting the results as a function of the number of exhaust particles per contrail flight meter made the figures much easier to interpret. We have therefore changed the x-axis in all panels of Fig. 5 accordingly and revised the related discussion in the manuscript. We have also added a figure showing EF_{pfm} as a function of particle radius, to still illustrate how the average EF_{pfm} changes for a given size distribution as a result of the parameters varied in our sensitivity analysis.

Minor comments:

1. In the abstract, it is not immediately clear to the reader what your base case assumptions are. I recommend stating them briefly.

Response: We think it would involve too much numerical details specifying the modeling assumptions in the abstract. We hope the reviewer and editor finds this position acceptable.

2. Line 16: The clouds themselves are not artificial; they are artificially produced.

Response: Fixed!

3. Line 37-38: What about NO_x emissions and the associated formation of nitric acid in hydrogen combustion cases? Could this contribute to volatile particle formation?

Response: We agree that the possible role of NO_x-derived species in particle formation from hydrogen combustion should be acknowledged. However, there is still considerable uncertainty regarding this pathway. To our knowledge, the formation of a distinct particle class solely from NO_x derivatives has not been demonstrated, while uptake of nitric acid by lubrication-oil droplets, and the resulting modification of their properties, remains a plausible but as yet unexplored possibility. We therefore do not include this pathway explicitly in our analysis, although parts of its potential effect may be indirectly encompassed by our sensitivity analysis of lubrication-oil characteristics, including size distribution, emission index, and hygroscopicity. We now mention NO_x in the Introduction (line 40) and make its omission explicit in Sect. 3.1.2 (lines 189–194).

4. Line 59-61: There is a dedicated simulation study on the potential role of lubrication oil particles in contrail formation (Zink et al., 2025).

Response: We thank the reviewer for pointing this out. The dedicated simulation study by (Zink et al., 2025) is now cited and discussed in the Introduction lines 63-66 and compared to our results in the Discussion lines 557-563.

5. Line 67-68: Wasn't it demonstrated by Ponsonby et al. (2024) that lubrication oil can act as condensation nuclei?

Response: It was, we were referring to the magnitude of hygroscopicity of oil, the sentence, now lines 70-72, has been modified for clarity.

6. Line 106: You use a value of 123 MJ kg⁻¹ for the specific combustion heat. Others use 120 MJ kg⁻¹ (Schumann, 1996, Bier et al., 2024). Could you clarify the source of this difference, and discuss its potential implications for your results?

Response: We thank the reviewer for pointing this out. As we understand it, the use of 120 MJ kg⁻¹ is conventional in much of the contrail literature, whereas values closer to 123 MJ kg⁻¹ are also used in some aircraft and propulsion studies. For example, (Khan et al., 2022) use 122.8 MJ kg⁻¹ as the mass-specific fuel energy content of hydrogen when converting between energy-based and mass-based emission indices. We originally adopted this higher value when implementing the model. A slightly higher value of the specific combustion heat could, in principle, lead to slightly higher plume supersaturation and therefore slightly stronger contrail effects. To assess the sensitivity of our results to this choice, we constructed an additional emulator using $Q = 120 \text{ MJ kg}^{-1}$ and ran CoCiP for cases with geometric mean radii of 5 and 25 nm, geometric widths of 1.28 and 2.72, and our standard set of 20 random dates over the latitude range $[-70, 70]$, longitude range $[-180, 180]$, and pressure range 200–300 hPa. For our original simulations using $Q = 123 \text{ MJ kg}^{-1}$, the average EF_{pfm} was 1.5% and 0.35% larger for the 5 nm cases with geometric widths 1.28 and 2.72, respectively. For the 25 nm cases, the corresponding differences were very small: the average EF_{pfm} with $Q = 123 \text{ MJ kg}^{-1}$ was $6.1 \times 10^{-3}\%$ smaller for geometric width 1.28 and $4.8 \times 10^{-4}\%$ larger for geometric width 2.72. These

differences are too small to affect any of our conclusions, and we therefore do not expect this choice to hinder comparison with previous studies.

7. Line 159-162 and line 229-230: Ponsonby et al. (2025) state that the K15 model is not capable of realistically resolving competition effects among different particle types. Please review this statement and discuss in more detail how this influence your results.

Response: We thank the reviewer for pointing this out. The limitation of the K15 model in resolving competition between different particle types is now discussed more clearly in Sect. 3.1.1 (lines 170–173), together with a brief explanation of how this may influence our results.

8. Line 232: Stating that the fit is “very good” is rather qualitative. Please quantify the errors between the original K15 model and your fit to provide a more precise assessment.

Response: We thank the reviewer for this important comment. Rather than relying on qualitative descriptions alone, we now quantify the disagreement through comparisons of EFp_{fm} between CoCiP using our emulator and CoCiP using the modified K15 model implemented directly. Selected values are now given in Sect. 3.1.4 (lines 265–269), while a more detailed assessment is presented in Supplement S6.

9. Figure 1: Your fits occasionally fail to reproduce the sharp decrease of activation fraction to zero at high ambient temperatures. Homogeneous freezing temperatures of supercooled droplets are typically below ~235 K, yet in some extreme cases, your fits show activation fractions above zero up to 250 K. Would it be beneficial to introduce a hard cut at 235 K?

Response: We thank the reviewer for this keen observation. A hard cut at 235 K could indeed be justified. However, these extreme cases are rare and only seldom satisfy the Schmidt–Appleman criterion. To assess the influence of this issue on our results, we performed an additional sensitivity test in which a hard cut at 235 K was imposed.

Relative to the original formulation, the hard cut changed the average EFp_{fm} by $-1.1 \times 10^{-4}\%$ for the geometric mean radius 5 nm, geometric width 1.28 case, -0.59% for 5 nm, 2.72, -1.3% for 25 nm, 2.72, and $+5.5 \times 10^{-6}\%$ for 25 nm, 1.28. These differences are sufficiently small that they do not affect our conclusions or the overall interpretation of the results.

10. Table 1: You provide the unit “nm” for the standard deviation, but it should be unitless, correct?

Response: We thank the reviewer for pointing this out. We have revised Table 1 to report this quantity as the geometric width, and the unit has been removed.

11. Table 1: Your oil emission index is based on measurements behind kerosene engines and is given in mg (kg fuel)^{-1} . Care must be taken when applying this emission index to a different fuel type. The lubrication system is separate, and oil emissions are not necessarily related to fuel flow. For example, if both a kerosene and a hydrogen engine emit the same oil mass from their lubrication systems, the emission indices will differ due to the different fuel flows, even though the actual emitted oil mass is the same. This may not strongly affect your conclusions, especially since you also varied the oil mass emission index, but it would be helpful to include a brief clarification on this point.

Response: We thank the reviewer for this important comment. We agree that applying an oil emission index expressed in mg (kg fuel)^{-1} to a different fuel type requires some care, since lubrication oil emissions originate from the lubrication system and are therefore not necessarily proportional to fuel flow. At present, we are not aware of any published estimates of lubrication-oil emission indices for hydrogen engines that would allow a more direct comparison. We therefore retain the present formulation, but now clarify this limitation explicitly in the manuscript (lines 283–287).

12. Figure 2 is based on the assumption that your model aircraft is present at each grid point at every timestep. Is that correct? In other words, no realistic air traffic is assumed, right? It is a valid approach to look at such potential contrail coverage. But a few more words on the used method would improve clarity.

Response: We thank the reviewer for this helpful suggestion. Yes, the model aircraft is assumed to be present at each grid point at every time step, so the analysis is based on a gridded potential-coverage approach rather than on realistic flight paths. We have now clarified this in the Introduction (line 82), the Models and Implementation section (lines 148–149), and the section describing CoCiP.

13. Line 425-428: CoCiP has been compared to measurement data for conventional kerosene combustion. It is not clear whether this validation also applies to scenarios with strongly reduced ice crystal numbers. Please clarify.

Response: We thank the reviewer for this important comment. We agree that the existing evaluation of CoCiP against measurement data for conventional kerosene combustion does not by itself establish its validity for cases with strongly reduced ice crystal numbers. To our knowledge, CoCiP has not been directly validated against observations in such a regime. It has, however, previously been used in an idealized sensitivity analysis of low-soot emissions (Rubin-Zuzic et al., 2025), although this does not constitute observational validation.

A key uncertainty in applying CoCiP to low-ice-crystal-number cases is its limited representation of the contribution from ambient aerosols and other volatile or ultrafine particles to ice particle formation. In the present study, we explicitly modify this part of the model. Another potential limitation is that low ice particle numbers imply larger individual ice particles and therefore faster sedimentation in CoCiP, owing to

its monodisperse particle-size assumption. It therefore remains uncertain how accurately CoCiP represents cases with strongly reduced ice crystal numbers relative to more conventional cases. We now discuss this explicitly in the Discussion (lines 533–536).

14. I appreciate the discussion about the comparison of CoCiP and APCEMM.

However, I am not happy about the sentence 'It is still uncertain if APCEMM or CoCiP is closer to the real world'. For any model, there is of course kind of uncertainty how close the results are to reality. But Akhtar Martínez et al. (2025) demonstrated that an important physical process is not well represented in CoCiP: the assumption of monodisperse ice particle sizes leads to the same sedimentation rate for all ice crystals, which produces a fallstreak and cuts off the lifecycle. As you note, APCEMM shows better agreement with large-eddy simulation studies.

Response: We thank the reviewer for this comment. We agree that the original sentence was not well phrased, as it did not reflect the specific physical limitation identified in CoCiP. We have therefore revised it to state more clearly that the monodisperse ice particle size assumption in CoCiP limits its representation of sedimentation and fallstreak development, whereas APCEMM, as a higher-fidelity model, shows better agreement with large-eddy simulation studies (lines 530-541).

15. This raises questions about the reliability of the radiative effect estimates in the present study. I recommend that this uncertainty be explicitly acknowledged, and I would also welcome a statement about it in the abstract.

Response: We thank the reviewer for this important comment. We agree that this uncertainty should be acknowledged explicitly, particularly in relation to the reliability of the radiative effect estimates. We have therefore strengthened the discussion of this point in the revised manuscript for example in the Discussion lines 545-552 and now also mention it in the Abstract lines 14-17.

16. Table S1: Should the unit of aerosol number concentration be cm^{-3} ?

Yes! Thanks!

17. S3: Do the relative humidity values $rh=0.1 \dots 1$ refer to relative humidity values over water? Only the combinations of T_a and rh are relevant, where the environment is supersaturated over ice (then persistent contrails develop). Would restricting to these values results in a better fit?

Response: We thank the reviewer for this relevant suggestion. The relative humidity values do refer to relative humidity over water. We chose the emulator range to cover all relative humidity values for which CoCiP Grid produced non-zero EFpfm output; in our simulations, the lowest such value was $rh = 0.11$. This reflects the fact that activation in CoCiP Grid occurs before downwash, whereas contrail persistence is evaluated after downwash.

Restricting the fit to ice-supersaturated conditions would therefore not necessarily be appropriate for the activation emulator itself. In addition, the main difficulty in fitting the emulator appears to arise less from the relative humidity range than from differences in shape across the size distributions considered. To test whether a restricted humidity range would improve the fit, we fitted an emulator using only cases with $rh > 0.6$ and evaluated it for two cases. This led to mixed results, reducing the underestimation from 4% to 0.8% for the 5 nm, 1.28 case, but increasing the overestimation from 6% to 7% for the 5 nm, 2.72 case, relative to direct implementation of K15 in CoCiP.

18. Figures S5 and S6: Could you combine them into a single figure, where you plot the difference between the two approaches?

Response: We thank the reviewer for this suggestion. Instead of combining Figs. S5 and S6 directly, we added a box-and-whisker plot summarizing the differences in EF_{pfm} across all contrails generated in simulations with four flight levels, at 60 grid points, and 20 days, comparing an implementation of K15 for soot in CoCiP with CoCiP's native soot activation function.

Technical Corrections:

- Line 71: It seems that something went wrong with the citations.
- Line 86: Just cite Kärcher et al. with \citet(Zink et al., 2026)
- Line 207: It seems that something went wrong with the citations.

Response: We thank the reviewer for pointing out these citation issues. We have corrected the malformed citations on lines 71 (now line 77) and 207 (now line 243), and revised the sentence and citation on line 86 (now lines 94-95).

Relevant changes made in manuscript

1. Inclusion of smaller geometric mean radii in the sensitivity analysis.
2. Inclusion of a sensitivity analysis of ambient aerosol number concentration.
3. Quantification of the differences between the emulator and the modified K15 model.
4. Quantification of the differences between CoCiP's native soot activation function and the modified K15 model.
5. Addition of quantitative statements throughout the manuscript, with corresponding revisions to the discussion of results.

6. Revision of the x-axis in the panels of Fig. 5, with corresponding revision of Sect. 4.2.
7. Addition of discussion of homogeneous droplet nucleation and NO_x-related particle-formation pathways.
8. Tempering of statements concerning engineering and technical implications.
9. Expansion of the discussion of model limitations, including limitations of both K15 and CoCiP in relevant regimes.
10. Clarification that the model used is CoCiP as implemented in pycontrails, together with addition of the relevant citation for the gridded CoCiP implementation.
11. Revision of the zonal plots such that altitude is shown on the vertical axis and latitude on the horizontal axis.
12. Addition of caveats to the abstract, including the need for further validation through more comprehensive modelling and future measurements.
13. Revision of terminology and notation for particle size distributions, including replacement of mean radius with geometric mean radius and standard deviation with geometric mean width, in line with field conventions.
14. Thorough language revision, citation correction, and technical proofreading throughout the manuscript.

- Bier, A., & Burkhardt, U. (2019). Variability in contrail ice nucleation and its dependence on soot number emissions. *Journal of Geophysical Research: Atmospheres*, 124(6), 3384-3400.
- Bier, A., Unterstrasser, S., Zink, J., Hillenbrand, D., Jurkat-Witschas, T., & Lottermoser, A. (2024). Contrail formation on ambient aerosol particles for aircraft with hydrogen combustion: a box model trajectory study. *Atmospheric Chemistry and Physics*, 24(4), 2319-2344.
- Khan, M. A. H., Brierley, J., Tait, K. N., Bullock, S., Shallcross, D. E., & Lowenberg, M. H. (2022). The emissions of water vapour and NO_x from modelled hydrogen-fuelled aircraft and the impact of NO_x reduction on climate compared with kerosene-fuelled aircraft. *Atmosphere*, 13(10), 1660.
- Rubin-Zuzic, M., Bugliaro, L., Marsing, A., Wang, Z., Voigt, C., Simson, C., Kaiser, S., & Ziegler, P. (2025). Reduced contrail radiative effect for fleets with low soot and water vapour emissions. *Atmospheric Environment: X*, 100353.
- Zink, J., Unterstrasser, S., & Burkhardt, U. (2026). Contrail formation for aircraft with hydrogen combustion–Part 1: A systematic microphysical investigation. *Atmospheric Chemistry and Physics*, 26(4), 3125-3143.
- Zink, J., Unterstrasser, S., & Jurkat-Witschas, T. (2025). On the potential role of lubrication oil particles in contrail formation for kerosene and hydrogen combustion. *Journal of Geophysical Research: Atmospheres*, 130(12), e2025JD043487.

