

Response to Reviewer 1

We would like to thank the anonymous Reviewer for their careful review. We found all the remarks to be very relevant and addressing the Reviewer's comments has improved the manuscript. Below is our response to the comments on a point-by-point basis. The following convention for text fonts is used:

- Review comments in black color
- Our answers in blue color
- *Pieces of text taken from the revised manuscript in red color and in italic font*

General impression and recommendation:

The authors use a flight planning tool in combination with numerical weather forecasts and a contrail prediction tool to study three concepts of contrail avoidance with increasing awareness of the risk involved. This is the risk of unintendedly damaging climate when the rerouting required for contrail avoidance leads to higher fuel consumption and emissions such that the increase of the CO₂ climate impact is larger than the climate benefit by contrail avoidance. The three concepts are: 1) planning a cost-climate optimal flight, that is, the cost-optimal contrail avoidance route and fly it; 2) as 1), but then use ensemble weather forecasts to estimate the bandwidth of the potential climate impacts and fly the cost optimal route only if there is a (positive) climate benefit in almost all ensemble members; 3) optimize the route in all ensemble members (i.e., plan N flights), and combine each of these routings with the remaining members of the ensemble, then select from those flights that have low risk the one with the largest predicted climate benefit.

To my feeling, this is a good strategy of research and the paper is a valuable contribution to the collection of ideas how to deal with the uncertainties caused by the contrail-avoidance vs. fuel consumption trade-off.

In the following I'll give some ideas for further improving the paper. All these are eventually **minor comments**.

We thank the Reviewer for their overall positive assessment of the paper. We answer their comments below.

1) My most important comment refers to the description of the "risk-optimal strategy" in Section 2. This description is not good and misleading. The reader gets a full understanding only in Section 6. What I understand from Section 2 is this: For each of N members of an ensemble forecast a climate cost-optimal route is computed. The uncertainty-based risk assessment mentioned in the previous par. ("risk-informed strategy") is then performed on each of the N flight trajectories (if this is correct, please state it clearly). A subgroup M

We thank the Reviewer for their comment, although it seems that part of the comment is missing. We tried to reformulate the corresponding paragraph in Section 2 as follows: *Here, we propose an alternative method and compute a cost climate-optimal route for each ensemble member of the weather forecast, providing a candidate trajectory for each ensemble member (Fig. 1, risk-optimised strategy). The risk of unintentionally damaging the*

climate is then estimated for each of these candidate trajectories. Those for which the risk is higher than a given threshold are ruled out. Amongst the remaining routes, the selected route is that with the highest average climate benefit. By choosing that route, we guarantee that the risk is lower than the given threshold while the predicted potential for climate benefit is maximum.

2) I would also like to suggest to the authors to consider the distribution/clustering of their "big hits". Admittedly, this is not central to this paper, but it is still an important topic that, to my view, has received too little attention in the past. I suggest that the data is screened for big hits, that is, extract the flights with the highest 2% of contrail warming and check where they appear spatially and temporally. I expect some clustering and this should make big contrail prediction easier and with lower risk. It could, however, lead to problems with airspace congestion if a big hit cluster is to be avoided completely. I believe you can here enhance the importance of your paper with little effort.

We agree with the Reviewer that investigating the spatial and temporal properties of the flights forming "big hits", defined in this study as contrail-prone situations that can be avoided with a high potential climate benefit, is of great interest to readers. A simple study of the distribution of strongly warming contrails (which are, admittedly, slightly different than our definition of big hits), is already available in Teoh et al. (2024), their Fig. 9. We think that the value of the analysis proposed by the Reviewer lies in an extensive study of such big hits in space and time, worldwide and for multiple years, which is out of the scope of the present manuscript. We leave this analysis for future work.

Teoh, R., Engberg, Z., Schumann, U., Voigt, C., Shapiro, M., Rohs, S., and Stettler, M. E. J.: Global aviation contrail climate effects from 2019 to 2021, *Atmospheric Chemistry and Physics*, 24, 6071–6093, <https://doi.org/10.5194/acp-24-6071-2024>, 2024.

3) A point for the discussion section is this: The study uses the NAFC as the model region. Do you think that the risk-informed and -optimal methods can be applied to more congested air spaces like Europe. How should the nearby presence of cirrus clouds and other contrails be treated?

This is indeed an important question. It is not yet clear whether the radiative impact of contrails formed by consecutive aircraft flying in the same region is additive or sub-additive. In congested airspaces, it might be more efficient not to reroute flights when several aircraft fly in the same region, effectively concentrating contrail formation in a given region, rather than reroute some but not all of the flights. We added a paragraph to the discussion section to mention this limitation: *Moving beyond individual contrail modelling, the impact of contrail-contrail and contrail-cirrus overlap on contrail avoidance strategies needs to be investigated, in particular in congested airspaces where contrails frequently form next to one another. As the additivity of their radiative impacts is not established, rerouting all or most of the flights in such regions might not always be the most effective option.*

4) The expression "to reject any risk" (twice in the Abstract and also later in the text) is surprising. This sounds as you would think of a significance level of zero (that is, probability of an error of first kind is zero). In a system with random elements this is not possible. In aviation this would imply to stay grounded.

Following the Reviewer's comment, we replaced "when choosing to reject any risk" by: *at the lowest risk tolerance level*. We also adapted the wording in the introduction and in Section 5.

5) Line 22: I suggest to replace "contribution" by ERF. Contrails have the largest ERF, but perhaps not the largest contribution (to anything).

Modified as suggested.

6) Line 24: replace "saturation" with "relative humidity". Saturation is 100% and does not exceed it.

Thanks for spotting that, it's corrected.

7) Section 2, description of the "risk-informed strategy". This was difficult to understand, probably because the additional calculation of the cost-optimal route is mentioned at the beginning. The reader must think that this is the important step, but it is not, I believe. I suggest to present this method as follows: We start as in the risk-unaware method with a climate-cost-optimal routing, then we use the uncertainties to estimate a risk. If the risk is too large, we calculate the usual cost-optimal route and fly it.

Indeed, this way of presenting things is much clearer. The text was adapted in that regard: *This strategy can be improved without disrupting operational flight planning processes too significantly by including in the workflow one additional step related to the risk of unintentionally damaging the climate (Fig. 1, risk-informed strategy). As in the risk-unaware strategy, a cost climate-optimal route is first calculated. From the calculation of contrail climate impact uncertainties, the risk of unintentionally damaging the climate is then estimated. If this risk is below a given threshold, fixed for example by the airline policy, the aircraft is rerouted and flies the cost climate-optimal route. Else, the aircraft is not rerouted and flies the usual cost-optimal route.*

8) Section 3.2, last par: Correction of the humidity field is necessary since NWP models often underestimate RHi in ISSRs. Perhaps you should note explicitly that the problem that ISSRs are often not predicted at their actual position cannot be fixed easily and that it needs more data for assimilation.

Following the Reviewer's comment, we've added the following sentence to the paragraph: *Models struggle to reproduce the humidity field, in particular the precise location of ISSRs, partly because the quantity of humidity data used in the data assimilation process is too low, among other reasons (Hofer et al., 2024). While waiting for additional observational data to improve the forecast quality, multiple studies have proposed a correction for the humidity field of the ERA5 reanalysis [...].*

9) Section 3.3, around line 207: It might be that the uncertainty is not weather-dependent, however, the efficacy factor itself is probably weather dependent. As the efficacy somehow measures the integrated effect of fast feedbacks over the course of weeks or months, it will certainly vary with the weather. I suggest to clarify this point. Otherwise, to simply assume a constant factor for this study is as a first step certainly in order.

Thanks for spotting this, this wording is indeed a mistake. The sentence was reformulated as follows: *However, while this factor depends on the synoptic weather situation (Verma and Burkhardt, 2026), estimating its value or uncertainty on a flight-by-flight basis is out of the scope of this study. Thus, we consider the climate efficacy of contrails to be the same for all flights.*

10) Section 4.1: I was surprised that flight A that exploits the jet stream needs more fuel than flight B that has strong headwinds, whereas both flights are quite similar in distance. Is this perhaps a consequence of the two different aircraft? Is the B767 more efficient than the B777?

This is indeed a consequence of different aircraft types for flights A and B, not particularly because of their efficiency but mostly because of their size and mass. The B767-300ER aircraft is much lighter (maximum take-off mass: 187 tons), while the B777-300ER is much heavier (maximum take-off mass: between 351 tons), thus consuming much more fuel.

11) Section 4.2, 2nd par: Is there a (tentative) explanation for this surprising finding, namely that flight A is sensitive to wind shear direction, while flight B is sensitive to particle emissions?

We extended the discussion to provide additional insights, but also state that the full understanding of the different roles of the parameters is out of the scope of the study: *Part of the difference in the role of the enhancement factor of nvPM emissions is due to the different emission index of nvPM for both flights, as expected from different engines. The aircraft flying flight A (resp. B) is assumed to be equipped with a GE90-115B engine (resp. CF6-80C2B6 engine) for the estimation of nvPM emissions, leading to an average nvPM emission index of about $2.8 \cdot 10^{14} \text{ kg}^{-1}$ (resp. $7.8 \cdot 10^{14} \text{ kg}^{-1}$). As the nvPM emission index is higher for flight B, the enhancement factor may have a stronger relative impact for flight B than for flight A. A detailed attribution of the differing sources of variability between the two flights is beyond the scope of this study, owing to the complexity and the strongly nonlinear nature of the processes represented in CoCiP.*

12) Section 4.2, around line 328: It is even worse, because the nominal estimates can be very different from the actual weather.

Indeed, however in our study we did not compare the nominal (nor the average) impact to actual impacts or simply reanalysed impacts. In this situation, we'd rather stick with a slightly weaker statement - the actual situation may lie in an outlier of the weather ensemble. Thus, we chose to keep the statement as it is.

13) Section 4.2, end: Now it becomes clear what you mean with "zero risk". It can be zero because you use a finite sample. Anyway, I suggest a slight rewording, say "negligible risk" and add in brackets where it is mentioned first (with zero climate damages in our finite sample).

Following our reply to comment #4, we have adapted the wording "to reject any risk" where appropriate. However, we have chosen to retain the "0% risk" terminology in the results section, as this value directly reflects the finite size of the sample used in our analysis rather than an absolute absence of risk.

14) Section 6, line 446-447: The sentence is easier to read if the last part "when adopting..." is shifted after "number of rerouted flights".

Modified as suggested.

15) Sect. 7: I agree that "the estimation of the climate benefit of reroutings must not be reduced to only one deterministic modelling configuration". I am not sure, but isn't this what the MRV-system is just doing to "determine" the contrail impact of single flights? Do you know more and can you comment on this? I think, your statement is an important recommendation.

To the best of our knowledge, this is indeed what is done in the MRV framework. We added a statement to make the recommendation clearer: *In particular, this recommendation extends to the monitoring, reporting and verification (MRV) framework recently introduced by the European Commission, in which the non-CO₂ effects of individual intra-EU flights are currently quantified using a single deterministic weather forecast and a single configuration of the contrail prediction model (European Commission, 2024).*

Response to Reviewer 2

We would like to thank the Reviewer for their careful review. We found the recommendations to be very relevant, making the readers more aware of the current state of the science. Below is our response to the comments on a point-by-point basis. The following convention for text fonts is used:

- Review comments in black color
- Our answers in blue color
- *Pieces of text taken from the revised manuscript in red color and in italic font*

Review

This paper introduces the concept of "risk-aware" contrail avoidance strategies. In such a strategy, a trajectory is optimized against a contrail forecast, and the resulting climate impact is re-evaluated by ensemble members to create a risk score that captures the likelihood that such a re-routing will result in a climate benefit or harm. Additionally, the authors present a "risk-optimized" approach, where multiple trajectories are created based on the differing ensemble members, allowing an operator to select the resulting trajectory that achieves the desired climate benefit at an appropriate risk level. The contributions in this manuscript are novel, appear scientifically and technically sound. An exploration of robust optimization strategies for contrail avoidance is currently limited in the literature on contrail mitigation, and thus I see this paper as a valuable and timely piece of work. In addition, the paper is very well written and organized. I recommend the paper be accepted, and only have very minor comments as noted below.

We thank the Reviewer for their overall positive assessment of the study. Replies to the comments are below.

- The measure of risk used in this work is related to the fraction of ensemble forecast members that, when used as input to a contrail model, cause climate harm rather than benefit. The authors attempt to also capture model uncertainty by also adjusting model parameters and measuring how frequently these adjustments result in climate damage. While this presents a valid metric, the authors do not present any evidence that this risk score is well calibrated. For example, seems like it should be possible for no forecast ensembles to predict a climate harm, but for a reanalysis to later show a climate harm. Calibrating this risk score ultimately seems like a very difficult task, because it is difficult for us to directly measure the climate impact of an individual contrail, and ultimately, the risk should be calibrated against observations rather than reanalysis products and models. For this reason, I find it acceptable for such a calibration to not be within the scope of this study. However, I suggest that the authors place a discussion of this important limitation within either Section 1 or 2 of this paper.

The Reviewer correctly points out that for our method to be valid, the actual climate benefit of rerouting should be reliably predicted. Such a condition can be verified using e.g., rank histograms. However, the scarcity of contrail impact observations hinders this verification process, which is left for future work. Following the Reviewer's recommendation, we added a

small discussion about this issue in Section 2: *Moreover, the effectiveness of the proposed strategies relies on the actual climate impact of the rerouted flight being reliably predicted during the planning process. Such a condition can be verified using e.g., rank histograms (Bröcker and Ben Bouallègue, 2020). Whether this condition is met cannot be verified at present, as direct observations of the climate impact of individual contrails remain scarce. Future work should focus on verifying that this condition is met.*

Note that a discussion of this limitation is already included in the conclusion (which was slightly reworked in the revised paper), thus we did not extend too much the new discussion in Section 2: “In particular, in order for our risk-aware decision-making to be valid, the actual climate benefit of a rerouting should be reliably predicted. This condition can be verified using for example rank histograms, that are commonly used to assess the reliability of ensemble forecasting systems (Bröcker and Ben Bouallègue, 2020). This is still an open question that needs to be addressed. We strongly advocate for additional research in evaluating and verifying CoCiP and similar models against observations, before they are used for operational contrail avoidance. Until then, a first step would be to assess whether the climate benefit estimated using reanalysed meteorological data falls within the estimated variability from ensemble weather forecasts, which will be the subject of future work.”

- On lines 315-316, the authors state that, for Flight B, the climate impact is highly sensitive to the parameter controlling the enhancement of nvPM emissions, but the same is not true for Flight A. The authors should report the engine types assumed for each of the two flights. It should be noted that recent work has shown that CoCiP is less sensitive to nvPM emission indices when accounting for the activation of vPM emissions, especially in newer lean-burn engines. Further, more recent versions of pyContrails have models of vPM activation as experimental parameters. I do not see it necessary for the authors to re-run their simulations with these experimental parameters, but I do recommend the authors include a reference to Ponsonby et al (<https://acp.copernicus.org/articles/25/18617/2025/>), and include a comment, possibly in Section 3.3, noting the above and that research in this space is evolving.

Following the Reviewer’s comment, we’ve added the following text to Section 3.3: *Recent work showed that in addition to nvPM emissions, ice crystal formation is dependent on volatile particulate matter (vPM) emissions, especially in new lean-burn engines (Ponsonby et al., 2025). While work is underway to include such findings in CoCiP, we chose not to include these experimental features in our study, but they could be included in the CoCiP parametric uncertainty estimation in future work.*

We also extended the discussion on the role of each parameter in the second paragraph of Section 4.2, explicitly stating the assumed engine types: *Part of the difference in the role of the enhancement factor of nvPM emissions is due to the different emission index of nvPM for both flights, as expected from different engines. The aircraft flying flight A (resp. B) is assumed to be equipped with a GE90-115B engine (resp. CF6-80C2B6 engine) for the estimation of nvPM emissions, leading to an average nvPM emission index of about $2.8 \cdot 10^{14} \text{ kg}^{-1}$ (resp. $7.8 \cdot 10^{14} \text{ kg}^{-1}$). As the nvPM emission index is higher for flight B, the enhancement factor may have a stronger relative impact for this flight. A detailed attribution of the differing sources of variability between the two flights is beyond the scope of this study, owing to the complexity and the strongly nonlinear nature of the processes represented in CoCiP.*

- In Section 6, the authors explore the concept of risk-optimized avoidance, where flights are optimized individually against a number of ensemble members. This concept is introduced in Section 2, where it is compared to by Simorgh et al, where risk is directly incorporated into the objective function of the optimization process. The authors' approach to risk optimization is valid and has the advantage of being far easier to implement. However, missing from both Section 2 and Section 6 is a comment on the limitation of the author's alternative approach. Namely, that by jointly optimizing against multiple ensemble members simultaneously, it may be possible to construct trajectories that achieve lower risk scores with similar operational costs. It is not clear a priori what advantage such a scheme would achieve, if any, and so this is a topic that would warrant further research in a future study.

We added a discussion in Section 2, where the risk-optimised strategy is described, to explain the main differences between the two approaches: *However, each candidate trajectory in our approach is optimised against a single ensemble member rather than jointly against the full ensemble. While our approach is simpler to implement, it may be possible to construct trajectories that achieve lower risk scores with similar operational costs using a single optimisation in which all ensemble members are considered simultaneously, as done by Simorgh et al. (2024).*

We emphasise that no matter the approach, the idea of the risk-optimised strategy is to process ensemble members directly in the flight planning system, rather than after an alternative route is made available.

The following comment is entirely editorial. The authors may consider this comment in the revision of this work at their discretion:

- Section 5 of your paper shows that the majority of contrail warming may be avoided through trajectory optimization at a relatively low risk level. This is somewhat a corollary of the main results of this paper. I have a concern that the presentation of the current paper may lead some readers to reach a different conclusion. Specifically, in Section 4, the authors show an example of a flight with a relatively high risk level. Based on the statistics in Section 5, it appears that this higher risk flight is somewhat of an outlier. These examples are still highly useful to illustrate the merit of the authors proposed risk mitigation strategy. To alleviate this concern, in Section 4, the authors may consider quantifying how likely it is to encounter flights like Flight A and Flight B. Further, the authors may consider switching the order of Sections 4 and 5, which would help provide more context to the reader for how typical Flights A and B are.

We thank the Reviewer for sharing this concern. Indeed, flights A and B are not randomly picked flights, but were chosen to nicely illustrate the objective of the risk-aware strategies. Following their recommendation, we added at the end of Section 4 a short analysis of the occurrence of flights A and B in the subset of flights: *However, flights A and B are not representative of the entire flight subset. For the 137 flights for which the nominal benefit is higher than 100 tCO₂e, the average estimated risk is 5%, with 71 reroutings for which the estimated risk is 0% and 9 for which the estimated risk is higher than 30%.*

However, we chose to remain as pedagogical as possible, going from the example to the general case, so we decided not to switch the order of Sections 4 and 5.