Travel Impact Model (TIM)
ADVISORY COMMITTEE

TECHNICAL BRIEF:

Communicating contrail impacts

January 2025

EXECUTIVE SUMMARY

This brief to the Google Travel Impact Model (TIM) Advisory Committee (AC) details the progress made on the topic of communicating contrail impacts in the TIM and was prepared in advance of the 5th AC meeting, held in June 2024. It summarizes the research, discussions, and development carried out by the Engineering team at Google, the International Council on Clean Transportation (ICCT), and the Task Group of experts that are a part of or delegated by the TIM AC.

This document presents the previously agreed upon climatological approach to communicate the estimated contrail warming impact of a future flight to the consumer at the time of booking. It elaborates on four different methods of communicating the impact, three of which rely on classifying flights into broad buckets to represent their impact, rather than presenting consumers with the estimated numerical value of the climate impact. The decision to classify, rather than present a number, was made at the 4th AC meeting, to acknowledge that any estimate of contrail impact for a flight in the future is highly uncertain because it is impossible to know the precise weather conditions that it will fly through. The three methods of communicating the impact via classification are assessed here on a variety of figures of merit, and case studies of specific routes are presented to aid the decision on which method to adopt.

At the 5th AC meeting, the decision was made to adopt Method 4, which classifies flights based on their estimated contrail impact relative to the CO_2 emissions typical of the route flown. Furthermore, it was suggested to consider incorporating seasonally varying probability of contrail formation rather than the constant value used in this analysis.

CONTENTS

GOALS

The goal of the workstream is to:

- Provide consumers with scientifically supported information about the non-carbon dioxide (CO_2) warming impacts of their flight choices
- Empower consumers to make the more sustainable choice when possible

While this workstream has worked on incorporating the warming impacts of all short-lived climate pollutants (SLCPs), this document focuses on the methodology for the inclusion of contrail impacts. It considers how contrail impacts are calculated for specific flights and investigates a few methods for communicating the impact to consumers. This focus on contrails is because they are expected to have the largest warming impact of the SLCPs.

Using a flat multiplier on CO_2 emissions to represent the global average contrail impact is not useful for consumers. It does not accurately represent the climate impact of individual flights, and because it does not offer additional differentiation between flights, it cannot steer consumer behavior toward a specific flight choice. We know there are average trends that can influence the contrail impact of a given flight and we chose to differentiate between the contrail impacts based on the origin, destination, local time of day, and season.

For the purposes of this document, we define "route" as an origin-destination pair and an "itinerary" as specific flight(s), direct or connecting, for a route. When searching for flights on an online booking platform, a consumer typically enters a route that they want to fly and then chooses between different itineraries for that route.

CLIMATOLOGICAL APPROACH

The Travel Impact Model (TIM) provides customers with emissions information at the time of booking a flight that is happening in the future. It is impossible to predict, due to uncertainty in weather forecasting, the contrail warming impact of an individual flight more than 24 hours in the future. Much of the uncertainty is associated with the science of predicting where and when a persistent contrail will form, which is still nascent.[1] However, we can look at historical weather conditions and aircraft trajectories and extract geospatial and temporal trends. These can be applied to future flight schedules to provide an estimate of the anticipated warming if a persistent contrail forms. That estimate is unlikely to be accurate for a specific flight, but it is directionally accurate when the impact of that flight is averaged over the course of a year or season.

1 Klaus Gierens et al., "How Well Can Persistent Contrails Be Predicted?" *Aerospace* 7, no. 12 (December 2020): 169. https://doi.org/10.3390/aerospace7120169.

One such climatological analysis was performed by Platt et al. and published in 2024.[2] They used the CoCiP contrail process model to simulate the energy forcing caused by a meter of contrail at each point on a 4-dimensional grid of latitude/longitude (every 0.5 degrees), altitude (3 pressure levels), and time of day (every 3 hours).[3] Each waypoint was run independent from the others, 100 times, for every fifth day of 2019. Each of the 100 runs used a randomly selected ERA5 weather ensemble to represent uncertainty in weather conditions. Each of 100 runs also drew key CoCiP parameters from representative distributions. The results of the simulations were aggregated to provide an estimate for the energy forcing per contrail kilometer (EFpcm) on a grid with points every 10 degrees latitude, 3 hours local time, and 3 months representing seasons. The result from the Platt et al. paper is shown below in Figure 1.

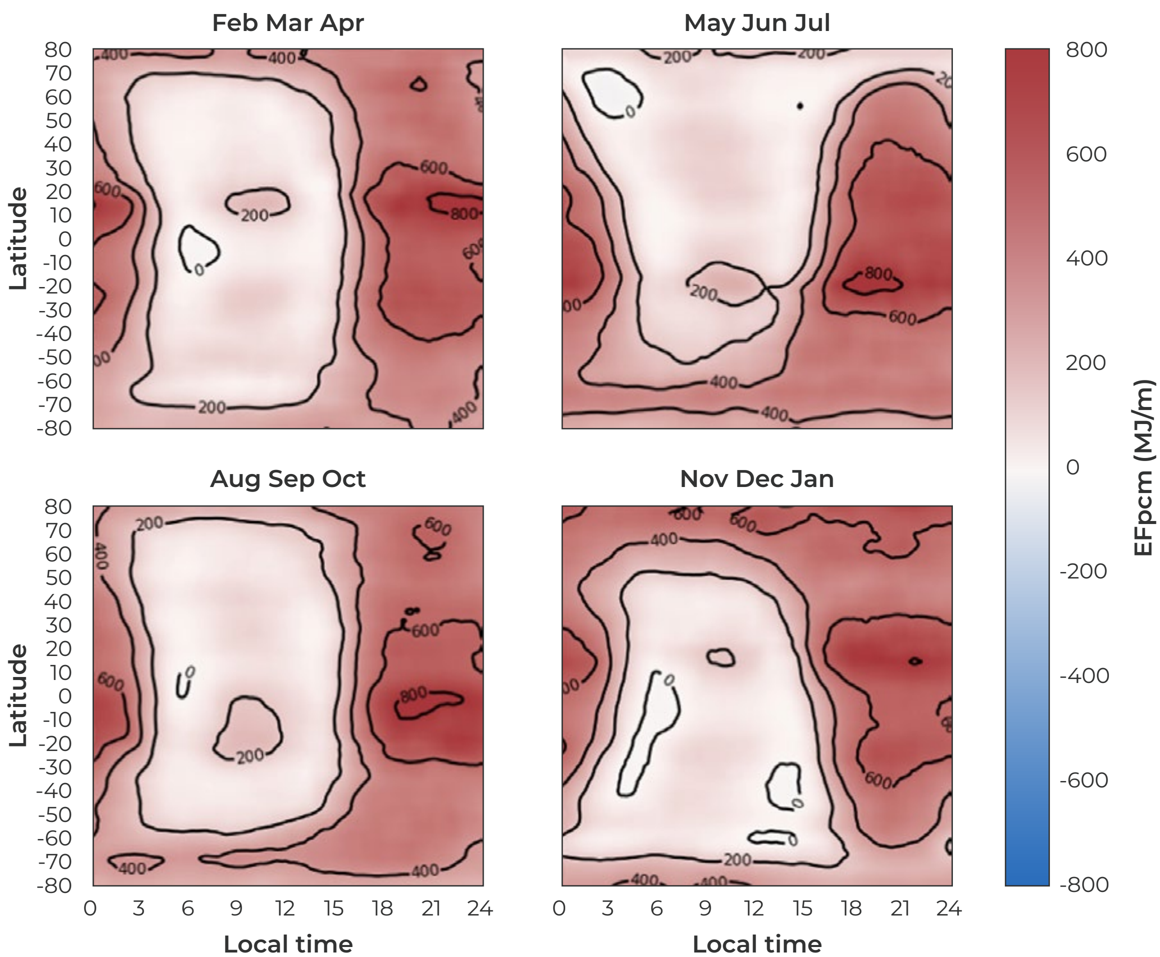

Figure 1. Mean EFpcm calculated for each 10 degrees of latitude, 3 hours of time of day, 3 months of the year for all flight levels and all parameter samples.

2 John C. Platt et al., "The Effect of Uncertainty in Humidity and Model Parameters on the Prediction of Contrail Energy Forcing," *Environmental Research Communications* 6, no. 9 (September 2024): 095015, https://doi.org/10.1088/2515-7620/ad6ee5.

3 Energy forcing (EF) refers to the integral of the radiative forcing of the contrail over its area and lifetime.

These heat maps indicate the EFpcm as a function of local time (x-axis), latitude (y-axis), and season (each one of the four maps). We discern a few trends.

- There is a diurnal pattern where contrails formed between 03:00 and 15:00 local time have a lower warming impact (indicated by the lighter colors) than those formed before 03:00 and after 15:00 local time.
- There is a seasonal pattern where the diurnal trend varies according to the daylight hours. For example, in May, June, and July, when the days are longer in the northern hemisphere, the lower impact area is larger. The same is true for the southern hemisphere in November, December, and January.
- There are small areas where contrails are expected to have a net cooling impact on average. These are the areas within the 0-level contour.

To convert an EFpcm for a latitude, time of day, and season to an estimate for the climate impact for a specific flight requires multiplying EFpcm values by the probability of persistent contrail formation and integrating it over the trajectory of the flight. Given the uncertainty in predicting persistent contrail forming on a future flight, a probability of 4.95% is assigned to all flights; future work will revisit this assumption to improve the accuracy of the TIM.[4] A great-circle distance trajectory is assumed between two airports and the scheduled local arrival and departure times are used to define the flight.

Examples of three itineraries overlaid on the heat map are shown in Figure 2. The first two itineraries are direct flights between JFK and SFO airports and the third is a one-stop flight from JFK to HNL. The direct flights are individual arrows that connect the airport latitudes over the duration of the flight. Even though the east-to-west flight (JFK to SFO, yellow arrow) is longer in duration, the west-to-east flight (SFO to JFK, green arrow) shows up as a longer arrow on the x-axis because the direction of travel is opposite to the rotation of the earth, which causes the flight to span a larger difference in local time. The one-stop itinerary is represented by two arrows, one for each flight and separated by the stopover time. Integrating the EFpcm values under these arrows over the idealized trajectory provides the EF for the whole flight. However, not all of the energy forcing contributes to surface temperature change. Atmospheric adjustments in the stratosphere and troposphere reduce the effectiveness of the energy forcing in changing the surface temperature.[5] To account for this, an efficacy factor of 0.42 is used to adjust the EF value to give an effective energy forcing (EEF) value.[6]

4 Roger Teoh et al., "Global Aviation Contrail Climate Effects from 2019 to 2021," *Atmospheric Chemistry and Physics* 24, no. 10 (May 2024): 6071–93, https://doi.org/10.5194/acp-24-6071-2024 .

5 James Hansen et al., "Efficacy of Climate Forcings," *Journal of Geophysical Research*, 110 (September 2005): D18104, https://doi.org/10.1029/2005JD005776.

6 D.S. Lee et al., "The Contribution of Global Aviation to Anthropogenic Climate Forcing for 2000 to 2018," *Atmospheric Environment*, 244 (January 2021): 117834, https://doi.org/10.1016/j.atmosenv.2020.117834.

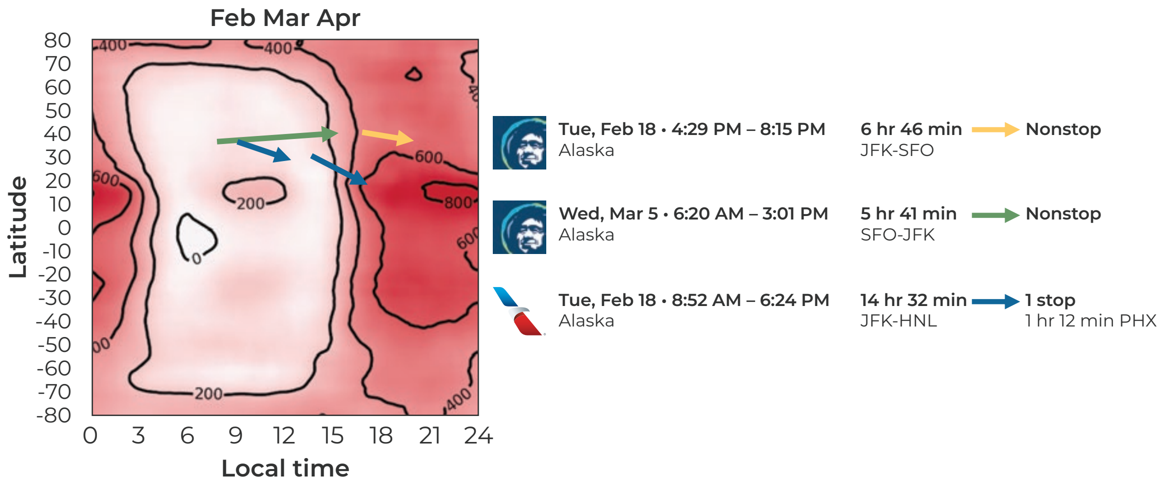

Figure 2. Three itineraries represented in the latitude and local time space.

WAYS TO COMPARE ITINERARIES

In the TIM's Task Group and Advisory Committee (AC) discussions, four possible approaches to compare itineraries were proposed.

1. **Method 1 (Route-based energy forcing multipliers):** Comparing itineraries for each route and presenting the percentage difference in climate impact compared with the itinerary with the lowest impact.

   While this presents a numerical value for EEF differences between flights, this method may exaggerate the importance of those differences. The range of EEF values for flights on one route can span many orders of magnitude, while still representing low climate impact. This can result in misleading information if percentage differences of EEF are presented to customers. Consequently, this method was abandoned at the 4th AC meeting (AC/4).

2. **Method 2 (Route-based energy forcing buckets):** Comparing itineraries for each route and bucketing them into four bins that range from the lowest to highest impact for the route.

   Like the above method, there is no way to compare across different routes. Every route would have itineraries in bucket 1 and 4, even if the difference in expected contrail impacts is minimal. This method provides actionable information to consumers, but it may overstate the differences between flights.

3. **Method 3 (Global energy forcing buckets):** Comparing itineraries across all global routes and bucketing them into four bins that range from the global lowest to the global highest impact.

   This method provides a way to compare across routes. However, because the bucket thresholds are determined by the global range of climate impact, there can be cases where all the itineraries within a certain route land in the same bucket. This might be of limited use to the consumers.

4. **Method 4 (CO_2 relative impact buckets):** Bucketing itineraries based on impact relative to CO_2 emissions.

   This method allows consumers to both compare itineraries within a specific route and compare impact across different routes. It also provides physical context for each bucket. For example, the warming bucket indicates that the climate impact of contrails is likely between 0.2x and 0.7x the CO_2 impact, while the cooling bucket indicates a cooling impact.

ASSUMPTIONS

- Aircraft that have cruise altitudes below 30,000 ft do not produce contrails. Cruising altitude is determined based on the EEA model's cruise altitude assumptions.
- Flight distances below 400 km (great-circle distance) do not produce contrail warming because the routes are too short for aircraft to reach cruise altitudes above 30,000 ft.
- Probability of persistent contrail formation = 4.95%
- We use GWP100, as used by the Intergovernmental Panel on Climate Change (IPCC), to normalize contrail forcing to CO_2 in Method 4.[7]
- AGWP100 CO_2 = 92.5×10^{-15} yr W m−2 kg CO_2
- Surface area of the earth, $S_{earth} = 5.101 \times 10^{14}$ m^2
- Efficacy factor: $\frac{\text{Effective Radiative Forcing}}{\text{Radiative Forcing}} = 0.42$

ANALYSIS OF THE CONTRAIL IMPACT DATA

The contrail impact calculation was performed for roughly 34 million flights scheduled between May 5, 2024 and July 5, 2025. Contrail impact was calculated for all scheduled flights, including those that would not reach altitudes where Ice Super Saturated Regions form.

Distribution of the warming impact

The warming impact output is in EEF and in units of gigajoules (GJ). Below are some summary statistics for the EEF output of the analyzed set of flights.

- Median = 10,336 GJ
- Minimum = -18,105 GJ
- Maximum = 535,555 GJ
- 0.27% of flights have negative EEF

7 The decision to use GWP100 is not an insignificant one and it has implications for the impact that is communicated using Method 4. This choice was made because GWP100 is used by policymakers and the IPCC to express the impacts of non-CO_2 emissions in CO_2-equivalent emissions.

Figure 3 shows the cumulative distribution of the contrail warming impact by cumulative contrail distance. On the x-axis is the fraction of contrail kilometers and the y-axis is the fraction of contrail warming. The results of the climatological model are the blue line. The gray dashed line represents what this plot would look like if we assumed a flat multiplier on CO_2 impact. This is equivalent to assuming that all flights create contrails, and that all contrails have the same climate impact. The green line represents the results of the same CoCiP model, but using reanalysis meteorological data which is only available after the flight occurs. This line extends above 1.0 and then starts to come back down to 1.0 at the end, due to the meteorological analysis predicting a percentage of contrail kilometers having a cooling effect. **This analysis is only possible after a flight has happened,** and it would not be possible to achieve this distribution when forecasting the contrail impact of flights. With the climatological analysis, there is a much smaller percentage of contrail kilometers (0.27%) that have a cooling effect. The climatological averaging makes the cooling impact nearly disappear, and this result has been seen in other climatological averaging.[8]

A way to read the graph is to look at 0.2 on the x-axis, representing 20% of contrail kilometers, and note where each line lands on the y-axis. The flat multiplier (gray dashed line) would suggest that 20% of contrail kilometers represent 20% of the contrail warming. The climatological approach (blue line) would suggest that 20% of the contrail kilometers represent 43% of the contrail warming. The meteorological model (green line) would suggest that 20% of contrail kilometers are responsible for 91% of the contrail warming. Thus, the climatological approach captures more of the uneven distribution of contrail impacts than the flat multiplier, but not nearly as much as the meteorological approach does.

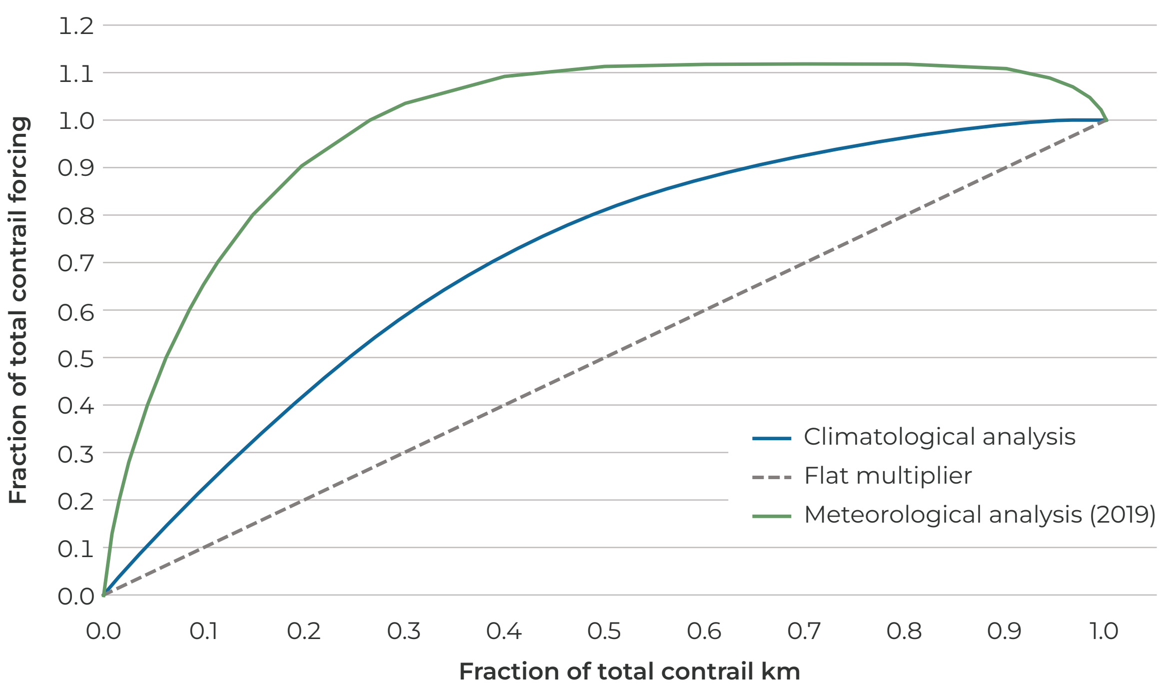

Figure 3. Distribution of the fraction of contrail kilometers and the warming that they cause. Flights are first sorted in descending order of contrail impact and then plotted on this figure.

8 Ines Sanz-Morère et al., "Impacts of Multi-Layer Overlap on Contrail Radiative Forcing," *Atmospheric Chemistry and Physics* 21, no. 3 (February 2021): 1649–81, https://doi.org/10.5194/acp-21-1649-2021.

We cannot expect the future predictions of the climatological model to reach the precision of the past-observation-based meteorological model. However, it provides a basis from which to distinguish between flights based on their likelihood of producing warming contrails and the magnitude of the warming effect of those contrails.

Figure 4 plots the frequency and cumulative distribution of the EEF of all of the approximately 34 million flights. For both graphs, the x-axis represents the EEF, expressed in GJ, and the y-axis represents the percentage of observations. This plot omits the flights with 0 or negative EEF values, to enable plotting with a log-scale x-axis. The EEF values range across 5 orders of magnitude, with the distribution centered around 10^4 GJ.

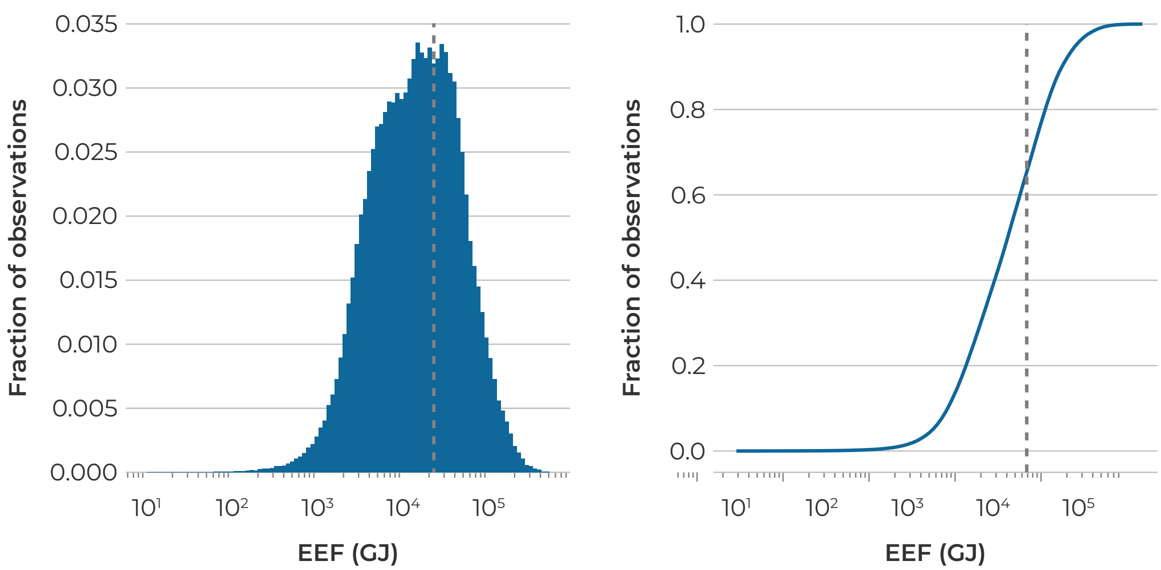

Figure 4. Frequency (left) and cumulative (right) distribution of the EEF values. Note the log-scale x-axis.

Distribution of the impact relative to CO_2 emissions

For the purposes of Method 4, we express the contrail impact relative to CO_2 impact of the flight. The non-CO_2 multiplier for a flight i on route j is defined as:

$$x_i = \frac{EF_{contrail,100,i} \times \frac{ERF}{RF}}{EF_{CO_2,100,j}}$$

$$EF_{CO_2,100,j} = AGWP_{CO_2,100} \times (365 \times 24 \times 60^2) \times S_{Earth} \times m_{CO_2,j}$$

Where $m_{CO_2,j}$ is the statistic representing the typical CO_2 emissions caused by flights on route j.

Summary stats for the multiplier x_i:

- Median = 0.34 (compared with Teoh et al.'s global average value of 0.32)
- Minimum = -0.31
- Maximum = 5.1

Figure 5 presents the frequency and cumulative distribution of the multiplier for all flights. For both graphs, the x-axis represents the non-CO_2 multiplier and the y-axis represents the percentage of flights.

The frequency distribution (left) shows two distinct modes, the first at ~0.2 and the second at ~1.2. There is also a long tail, with a maximum value of 5.1. In the cumulative distribution graph (right), the flights with zero contrail impact are represented in the 17% jump across the value of zero. Ninety-nine percent of flights have a multiplier that is less than 2.2.

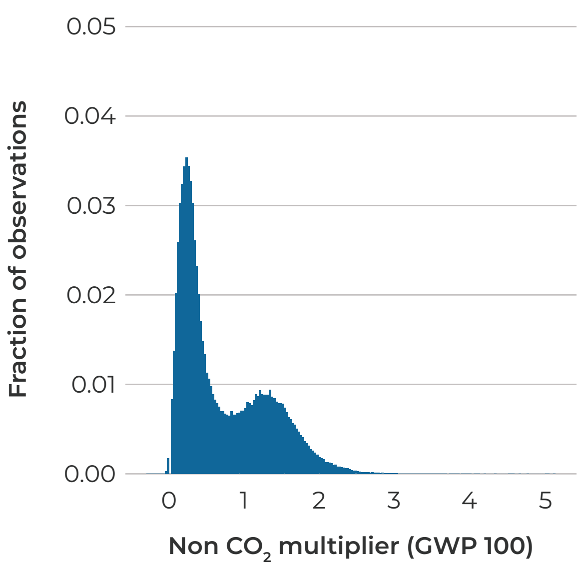

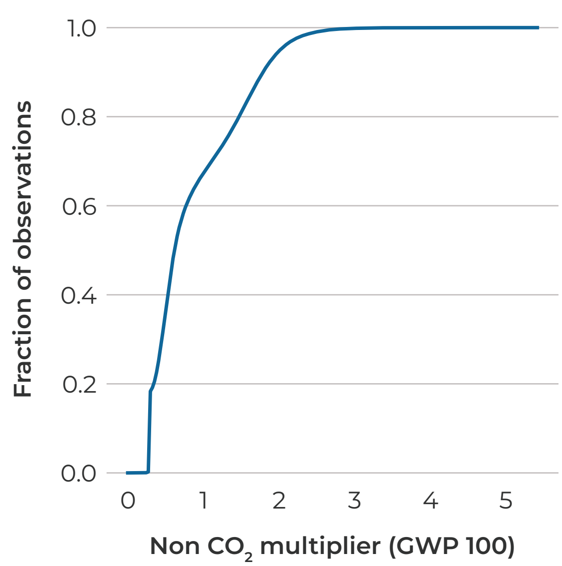

Figure 5. The frequency (left) and cumulative (right) distribution of the CO_2 multiplier across all flights.

METHODS OF COMMUNICATING IMPACT

In the TG and AC discussions over the past year, it was decided that numerical values for the contrail impact of a flight in the future will not be used because there is significant uncertainty in the numerical value. Instead, we decided to classify the impact into different buckets. The definition of the bucket depends on the method chosen. While we are not expected to define the bucket ranges immediately, for the purposes of this exploration, the following example bucket thresholds are defined.

- Method 2: Bucket ranges are chosen based on the range of values for an origin-destination-season combination. The buckets are equal width and range from the minimum to the maximum contrail impact value for a given combination.

- Method 3: The 25th, 50th, and 75th percentile of the EEF values, referred to as EEF_i, are used to define the four buckets. This places equal numbers of flights in each bucket.
  - Bucket 1: EEF_i < 2,778 GJ
  - Bucket 2: 2,778 GJ ≤ EEF_i < 10,336 GJ
  - Bucket 3: 10,336 GJ ≤ EEF_i < 29,200 GJ
  - Bucket 4: 29,200 GJ ≤ EEF_i
- Method 4: The bucket thresholds are defined to place equal numbers of flights in the Negligible, Warming, and Strongly warming buckets.[9]
  - Cooling (x_i < -0.2)
  - Negligible (-0.2 ≤ x_i < 0.2)
  - Warming (0.2 ≤ x_i < 0.7)
  - Strongly warming (0.7 ≤ x_i)

We can clearly communicate the thresholds of the buckets for Methods 3 and 4, and these threshold definitions are consistent for all routes. Communicating the thresholds for Method 2 would require publishing a list of all origin-destination-season combinations (of which there are 186,915 unique combinations) and the thresholds used for each combination. As this is more complex than a single definition of thresholds for all routes, **Methods 3 and 4 are considered more transparent than Method 2.**

Figure 6 represents the bucket definitions for Method 3 (left) and Method 4 (right). The x-axis represents the fraction of flights, and the y-axis represents the metric used (EEF or multiplier). The blue line represents the metric's cumulative distribution over the flights.

The equal width of the buckets indicates that the thresholds are designed to distribute equal numbers of flights in each bucket. In Method 3, each of the four buckets contains 25% of the flights. In Method 3, the Negligible, Warming, and Strongly warming buckets contain roughly one-third of the flights. The Cooling bucket only contains 0.001% of the flights and is thus not visible on this plot. Cooling flights are explored in more detail in the "Cooling flights" section at the end of this document.

---

9 Based on AC feedback, the "negligible impact" bucket is set to extend into the negative warming space, too. The threshold to be considered Strongly warming is set to 0.7. This means a flight that creates a contrail impact that is more than 70% of the CO_2 impact would be classified as a Strongly warming flight. A value of 0.7 is also roughly 2x the median contrail impact of a flight, according to our analysis. Teoh et al. calculated the global average impact to be 0.32x of CO_2.

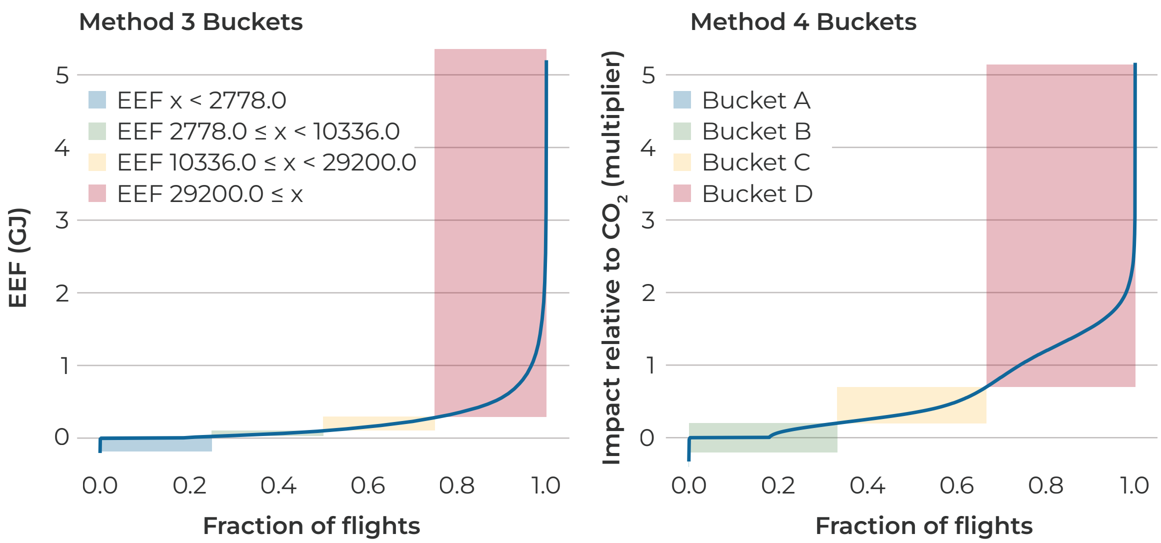

Figure 6. The bucket definitions for Method 3 and Method 4.

A similar graph cannot be produced for Method 2 because the bucket classifications happen individually for each origin-destination-season combination. However, we can calculate the percentage of flights in each bucket. The percentage of flights in each bucket for each method is shown in Table 1. The thresholds for Methods 3 and 4 are designed to ensure even distribution in the buckets. Method 2, however, is more polarized, with higher percentages in the lowest and highest bucket and only 13% of flights falling in the middle B and C buckets. This is an artifact of the classification in Method 2, which buckets the routes based on the range of EEF values for that route. If there are flights for the origin-destination-season combination with two or more contrail impact values, there will always be a Bucket A and a Bucket D.

Table 1. Percentage of flights in each bucket for each classification method

Percentage of flights in:	Route-based energy forcing buckets (Method 2)	Global energy forcing buckets (Method 3)	CO_2 relative impact buckets (Method 4)
Bucket A	56%	25%	0%
Bucket B	7%	25%	33%
Bucket C	7%	25%	33%
Bucket D	30%	25%	34%

Another way to look at the data would be to quantify the total warming represented by the flights in each bucket. Table 2 lists the percentage of contrail warming that is represented by each bucket, for each method. Ideally, we would like Buckets A–D to represent a monotonically increasing percentage of contrail warming. In other words, the total contrail impact of all flights classified in Bucket D should be greater than the total contrail impact of all flights classified as Bucket C, and so on. This increases the degree to which a method provides impactful choice (i.e., choosing a flight in a bucket with less warming reduces contrail impacts). We achieve this for Methods 3 and 4, with Method 3 being more polarized than Method 4. Method 2, on the other hand, would label flights that represent 21% of the total warming in Bucket A. This difference is driven largely by long-haul flights; these may have the least warming impact within a route, and hence are categorized in Bucket A, but create a high magnitude of contrail impact due to the length of the flight. With this, we can say that **Methods 3 and 4 have a better distribution of impact over the defined buckets than Method 2**.

Table 2. Percentage of contrail warming represented in each bucket for each classification method

Percentage of warming in:	Route-based energy forcing buckets (Method 2)	Global energy forcing buckets (Method 3)	CO_2 relative impact buckets (Method 4)
Bucket A	21%	1%	0%
Bucket B	8%	7%	4%
Bucket C	10%	20%	33%
Bucket D	61%	72%	63%

CONSUMER CHOICE

We would like to provide the customer with the ability to make a more sustainable choice wherever possible. To analyze how each method impacts what a consumer sees, we group flights into origin-destination-season combinations. The 34 million flights get grouped to form 186,915 unique origin-destination-season combinations.

Further, we filter for combinations that have more than 92 flights. This indicates more than one flight per day on average for that combination in a season (which lasts 3 months). This reduces the combinations down to 78,094 origin-destination-season combinations that have more than one flight per day. The following analysis is performed on these 78,094 combinations.

When the consumer has no choice

The consumer has **no choice** when all flights in an origin-destination-season combination are classified into just one bucket. A consumer has **at least two choices** when the flights in an origin-destination-season combination are distributed among two or more buckets. Figure 7 presents these percentages for each method as a pie chart, and the two percentages add up to 100%.

Method 2 provides the most consumer choice, with only 17% of the combinations having no choice and the remaining 83% having some choice (i.e., flights represented in two or more buckets). Method 4 provides choice in 72% of the combinations, which is marginally higher than for Method 3 (70%).

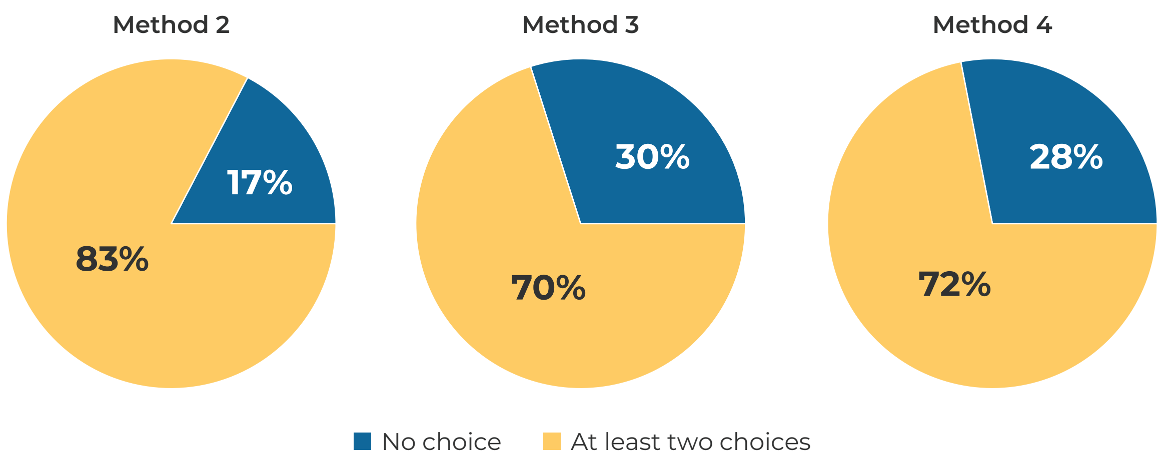

Figure 7. Percentage of origin-destination-season combinations that present the consumer with no choice or some choice.

Good versus bad choice

A consumer is said to have a **good choice** if at least one flight in the origin-destination-season combination is classified in Bucket A or B. A consumer is said to have only **bad choices** if all flights are classified in either Bucket C or D, but options in both buckets are present. A consumer is said to have **no choice, only bad options** if all flights for an origin-destination-season combination fall in one bucket and that bucket is either Bucket C or D.

Figure 8 presents these categorizations of choice for each of the methods. By design, Method 2 will always have at least one flight in Bucket A. Consequently, it provides a good choice for 100% of the origin-destination-season combinations. Method 3 provides a good choice for 77% of flights, with the remaining split between 12% with bad choices and 11% with no choice, only bad options. Method 4 provides a good choice for the fewest combinations, at 62%. However, the remaining split as 28% with bad choices and 10% with no choice, only bad options. While Method 4 has fewer combinations with good choices than Method 3, it has fewer cases of the no choice, only bad options scenario than Method 3.

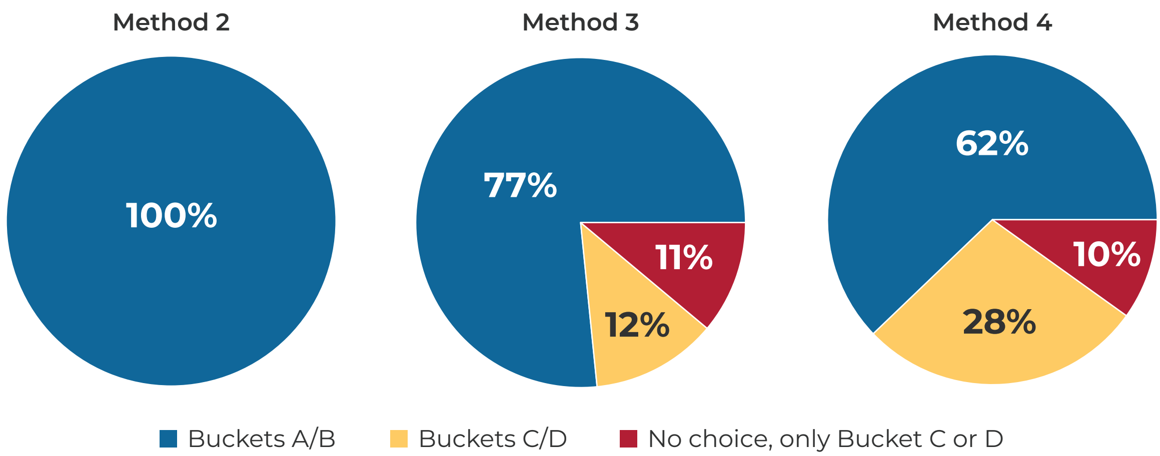

Figure 8. Categories of choices provided by each of the methods.

Note that these definitions of "good" and "bad" choices may not align with choices being impactful. Indeed, Methods 3 and 4, which arrange flights such that each bucket accounts for monotonically increasing share of contrail impacts, provide more impactful choices than Method 2.

However, if we are looking purely at the availability of choice, **Method 2 provides more choice to the consumer than Methods 3 and 4. Method 4 provides marginally more choice than Method 3.**

Case studies of FRA-LHR and LHR-JFK

To further investigate this question of consumer choice, we present the case of two routes in the summer season, Frankfurt to London (FRA-LHR) and London to New York City (LHR-JFK). FRA-LHR is an intra-Europe, short-haul route (657 km) and LHR-JFK is a transatlantic, long-haul route (5,554 km). The FRA-LHR-Summer combination has 1,385 flights, and the LHR-JFK-Summer combination has 2,120 flights.

Figure 9 plots the FRA-LHR-Summer results on the left and the LHR-JFK-Summer results on the right. Each row represents the bucket classifications for a specific method. All six plots show the distribution of flights and their contrail impact. The x-axis is always the percentage of flights. On four of the panels, the y-axis represents the EEF for Methods 2 and 3 and on the other two panels, it represents the non-CO_2 multiplier for Method 4. Note that the impact from the LHR-JFK flight is roughly an order of magnitude higher than the impact from the FRA-LHR flights for Methods 2 and 3. The blue line represents east-to-west direction flights, while the brown line represents the west-to-east flights. The background colors represent the buckets for each method.

First, let us ignore the bucket colors and only look at the variation in east-to-west (blue line) versus west-to-east (brown line) flights in the middle row of Figure 9. Shorter flights have less variation between the two legs of a round trip and there is more variation between the two legs for longer flights. The longer flights are traveling across more time zones, with east-to-west flights (blue lines) flying in the direction of the sun's movement. This means there is less local time variation in east-to-west flights. So, a flight taking off at night is more likely to fly entirely at night. This leads to a more polarized impact (flights with low impact are mostly day flights and flights with high impact are mostly night flights). West-to-east flights (brown lines) are flying against the movement of the sun and so there is a larger variation in the local time. A flight taking off during the day at noon might land in the evening, after the sun has set. This leads to a less polarized, more even distribution of impact.

Now, let us look at the bucket colors and the classification of flights in the buckets. First, by defining energy forcing buckets at the route level, Method 2 (top row of Figure 9) provides the greatest choice across both short-haul (left) and long-haul (right) flights. In contrast, Method 3 (middle row), which defines those buckets via all flights globally, categorizes most short-haul flights in the lower two buckets and most long-haul flights in the highest bucket. Method 4 (at bottom in Figure 9), by normalizing the contrail energy forcing by CO_2, is less sensitive to flight distance than Method 3 but more sensitive than Method 2.

Second, the east-to-west and west-to-east variation creates interesting differences in the classifications. Although there is little variation in the scoring of short-haul flights based on direction of flight, the scoring of long-haul flights on Methods 3 and 4 is sensitive to flight direction. Both methods assign more flights to the two highest warming buckets for the eastbound flight (JFK-LHR) than the reverse. Presumably, this is because all eastbound flights have some portion of the trip at night (more warming), while many westbound trips could occur completely during the day due to westward travel through time zones. You will notice that the Method 2 plots only include one direction of flight. This is because the bucket definitions change with flight direction and that is not legible on a single plot. However, as the bucket definitions are direction-dependent, the distribution of classifications does not change significantly.

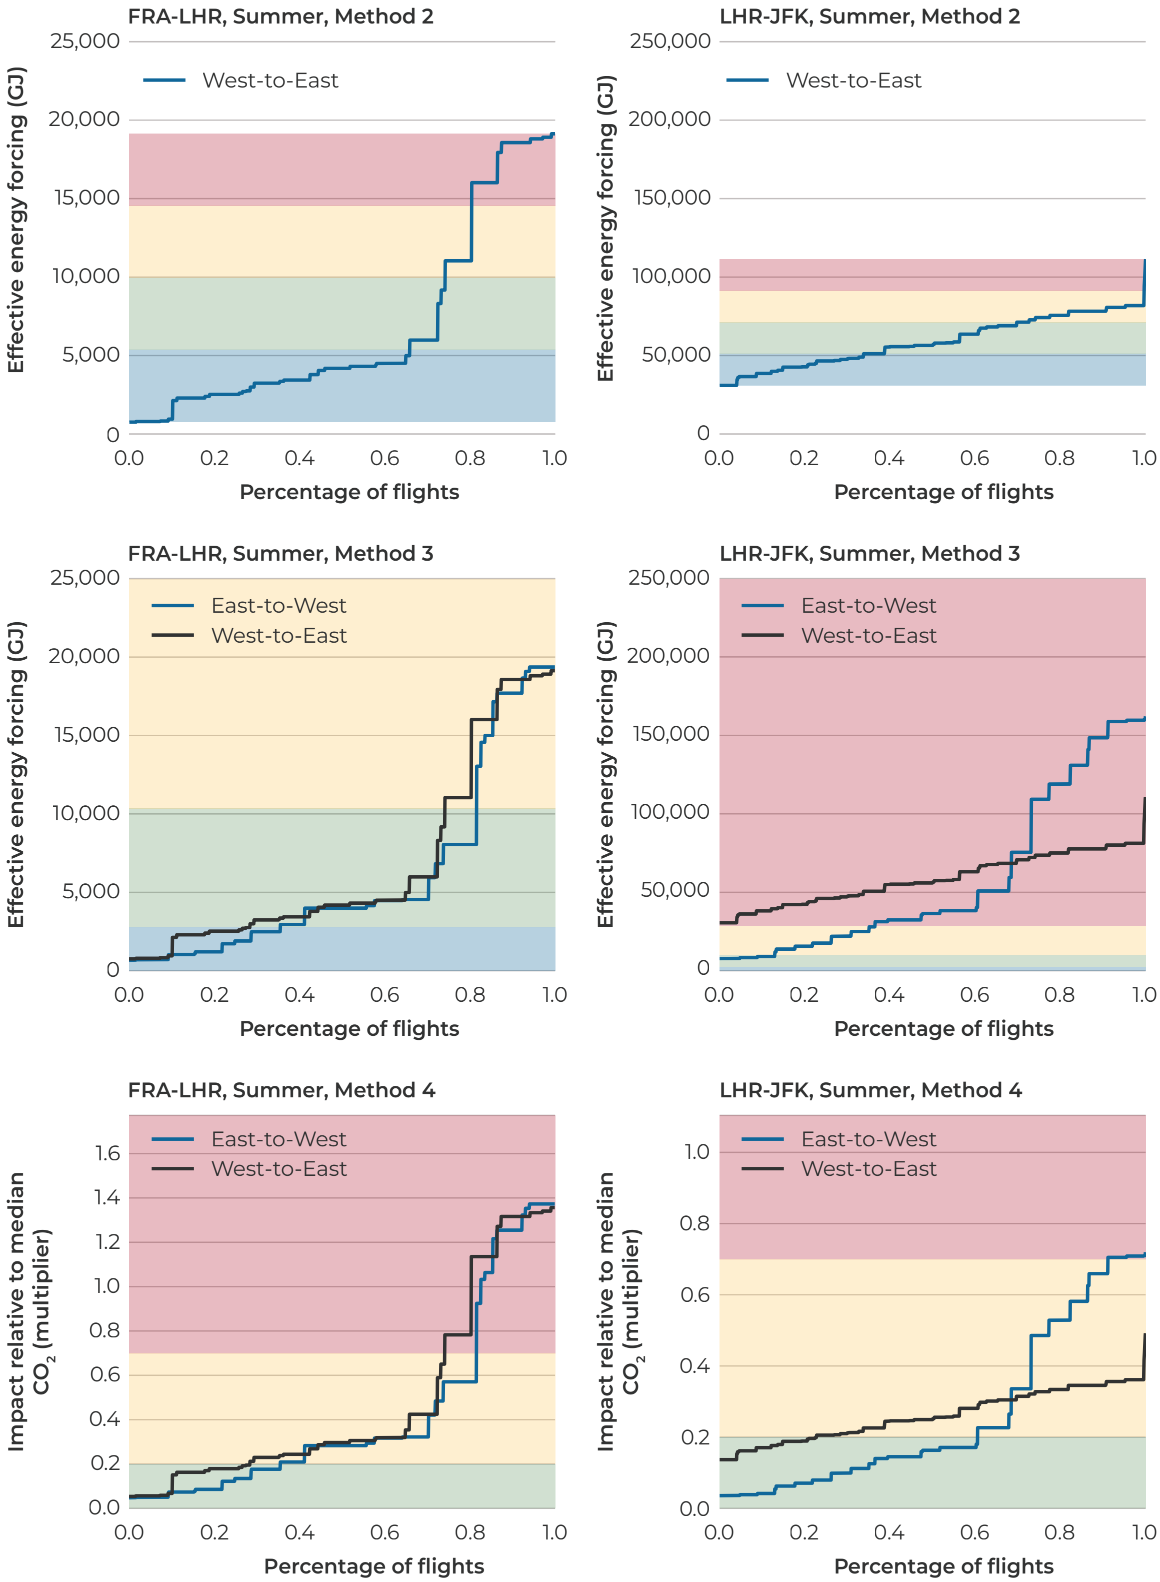

Figure 9. Case study in choice.

DISTANCE-BASED ANALYSIS

We would also like to see how the classification changes with distance. For this, we define four distance bands:

- Commuter (< 500 km)
- Regional (≥ 500 km, < 1,500 km)
- Medium-haul (≥ 1,500 km, < 4,000 km)
- Long-haul (> 4,000 km)

For each distance band, and for each method, we calculate the percentage of flights that fall into each bucket. Figure 10 presents this as a stacked bar graph, where the y-axis represents the percentage of global flights. The axis is divided into four parts, one for each distance band, as indicated by the label. The label also contains the definition of the distance band and the percentage of global flights within the distance band. Within each distance band, there are three stacked bars, one for each method. The total height of each stack within a distance band is the same and it represents the percentage of global flights in that distance band (the same value represented in the label). The division of flights into buckets is represented by the different colors.

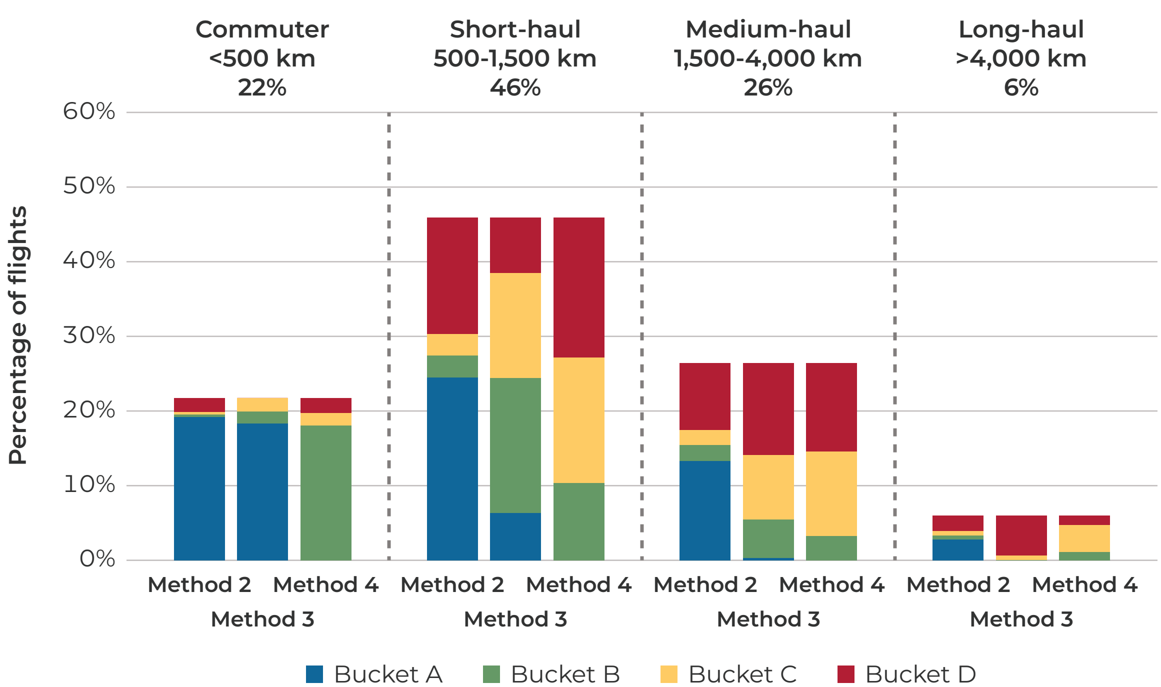

Figure 10. Percentage of flights in each warming bucket for Methods 2-4, for each distance band.

The overarching trend indicated by the total heights of the stacked bars is that the Short-haul segment accounts for the largest percentage (46%) of the flights and the Long-haul segment accounts for the smallest percentage (6%). There is a clear trend for Method 3 as it transitions from having most flights in Buckets A and B for the

Commuter and Short-haul distances, to having most flights in Buckets C and D for the Medium- and Long-haul distances. Method 2 has a consistent trend of classifying most flights as either Bucket A or Bucket D, regardless of the distance band. Method 4, due to its expression of contrail warming as a function of CO_2 emissions, has a more complex trend: It classifies greater percentages of Short- and Medium-haul flights in Bucket D than Long-haul flights. The greater CO_2 emissions of Long-haul flights means that even though the contrail warming impact may be greater in magnitude for Long-haul flights, the multiplier on the CO_2 emissions used to represent contrail warming for these Long-haul flights tends to be smaller.

Figure 11 is similar to Figure 10, except the y-axis is the percentage of contrail warming. As a result, the total height of each stack within a distance band represents the percentage of contrail warming caused by flights in that distance band (the same value represented in the label). The main change is that Medium-haul and Long-haul flights represent greater shares of the contrail warming than their flight frequency would suggest. This is to be expected, as the contrail warming impact of a flight is correlated with the length of the flight. The Commuter distance band only contributes 2% of the total contrail warming, even though it represents 20% of all scheduled flights.

Within the methods, similar trends persist. Classifications from Method 3 shift from being in the lower buckets to the higher buckets as the distance increases. Method 2 continues to classify most of the contrail warming as either Bucket A or Bucket D, regardless of the distance band. As before, Method 4 classifies most of the contrail warming for Short- and Medium-haul distances within Bucket D, but classifies most of the warming from Long-haul flights in Bucket C.

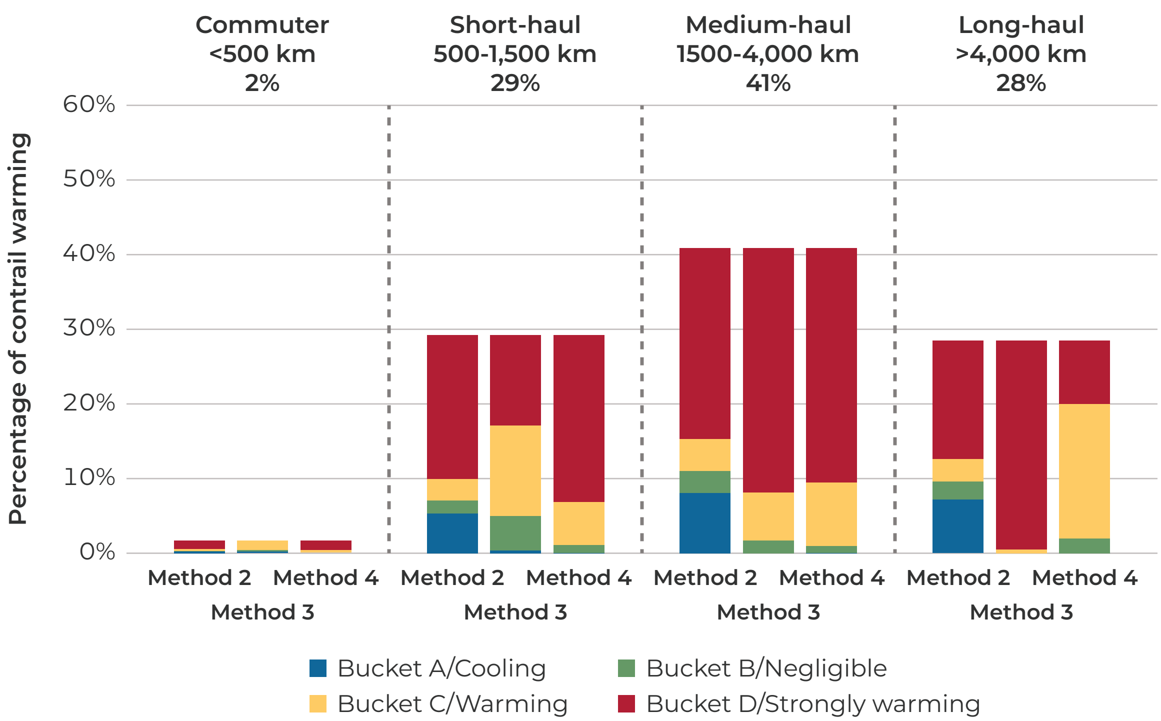

Figure 11. Percentage of contrail warming in each warming bucket for Methods 2-4, for each distance band.

Figure 12 is similar to Figure 11, except the y-axis is the percentage of **total warming** (CO_2 + contrail). Now the total height of each stack within a distance band represents the percentage of total warming caused by flights in that distance band and there is a monotonically increasing trend in the total warming associated with each distance band. Across Figures 10–12, we see that the 6% of Long-haul flights are responsible for 28% of the contrail warming and 36% of the total warming. In contrast, Short-haul flights, which are 46% of all departures, are responsible for 29% of contrail warming and 26% of total warming.

The same trends as in Figures 10 and 11 are seen within the methods. Classifications from Method 3 shift from being in the lower buckets to the higher buckets as the distance increases. Method 2 continues to classify most of the contrail warming as either Bucket A or Bucket D, regardless of the distance band. As before, Method 4 classifies most of the total warming for Short- and Medium-haul distances within Bucket D, but classifies most warming from Long-haul flights in Bucket C.

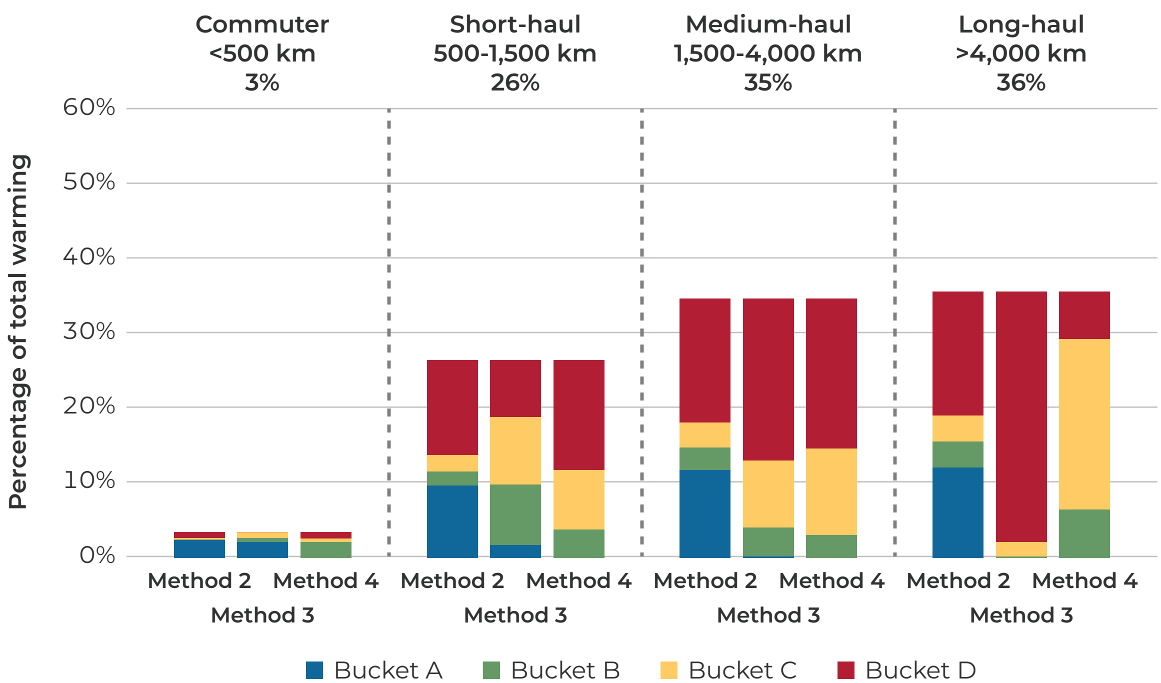

Figure 12. Percentage of total warming (contrail + CO_2) in each warming bucket for Methods 2–4, for each distance band.

Connecting flights

The analysis so far has focused on direct flights between two airports. The contrail impact of connecting itineraries can be determined by summing the contributions from individual flights. However, classifying them requires different treatment under the different methods.

For Method 2, with its route-based EEF buckets, this would be the most difficult to implement. It would require defining maximum and minimum EEF values across all possible itineraries. Defining the maximum EEF value for a route is difficult. As an extreme example, it is theoretically possible to fly between Frankfurt and London

Heathrow by connecting through Istanbul; even though this option on Turkish Airlines shows up on Google Flights, it would be absurd to include the EEF of that itinerary as the top of the range in defining the buckets for the FRA-LHR route.

Doing this for Method 3, with global energy forcing buckets, is easier. It would require a similar definition of the global maximum EEF, but this could continue to be defined as it has been in this analysis already, by choosing the maximum value across all direct flights. This could result in some connecting itineraries exceeding this maximum value, but those could be classified in the highest warming bucket without impacting the ability to communicate their impact.

For Method 4, with CO_2 relative impact, it is similarly easy, with no change in processes required. The sum of the EEF over the connecting flights would be divided by the average CO_2 emissions for the route. This average emissions value for a route, referred to as Market Reference Emissions, is already calculated for Google Flights and would not require any special treatment.

Integrating non-contrail SLCPs

An outstanding topic that has not been addressed in this document but has been discussed by the AC is the integration of non-contrail SLCPs, notably cruise NO_x, into the output. It is essential that the method chosen to communicate contrail impact is compatible with the communication of the warming impact of non-contrail SLCPs. Lee et al.[10] presented the impact of non-contrail SLCPs as both, as GWP100 and as average global radiative forcing per unit emission. Thus, they could be represented in both CO_2-equivalents and EEF. As such, all three methods could integrate the impact of non-contrail SLCPs.

Understanding the metrics

A question that was not settled at the Task Group meeting is: Which is the most understandable way of communicating the impact? Methods 2 and 3 are using EEF to express the warming, while Method 4 is converting the EEF into CO_2-equivalents and then normalizing them by the average CO_2 emissions for the route to express the warming as a multiplier of the average CO_2 emissions of the route. In all cases, the metric value is being abstracted away by classifying the warming values into buckets. To help make decisions, here are some arguments on both sides.

1. EEF is a more direct way to express the warming that is being caused by contrails. Expressing them relative to CO_2 requires a decision on a metric (in our case, we have chosen the GWP100).

2. EEF is an entirely new concept for consumers and might require a lot of education. Consumers are used to seeing the CO_2 impact of their flights. Defining the buckets based on multipliers on CO_2 impact could be more understandable.

10 Lee et al., "The Contribution of Global Aviation to Anthropogenic Climate Forcing."

3. Defining buckets based on multipliers would require consumers to do math to fully grasp the impact of their flights. With buckets defined on a global EEF range, Method 3 is perhaps simpler: Flights in the higher buckets are always contributing significant warming to the environment.

FINDINGS

Based on the analysis above, we can start to assess the methods on their fitness to the purpose of incorporating contrail impact into TIM outputs. We use a grading system similar to the one used for the alternative model selection. There is a scoring scale of -1, 0, or +1 that is represented by red, yellow, and green colors, respectively. Table 3 presents the grades we assign to each method along six figures of merit: transparency, distribution of impact, consumer choice, handling connecting flights, ability to integrate non-contrail SLCPs, and understandability.

Table 3. Grading the methods

Figure of merit	Method		
	Route-based energy forcing buckets (2)	Global energy forcing buckets (3)	CO_2 relative impact buckets (4)
Transparency			
Distribution of impact			
Consumer choice			
Connecting flights			
Other SLCPs			
Understandability			
Total	0	3	4

Because we can easily publish the bucket thresholds for Methods 3 and 4, they get +1 scores. Method 2 gets a 0 as we could publish bucket thresholds, but it would not be in an easily understandable form.

As the total warming associated with each bucket is monotonically increasing for Methods 3 and 4, they get +1 scores for distribution of impact while Method 2 gets a -1.

Method 2 provides the highest consumer choice and gets a score of +1. Method 4 provides marginally more choice than Method 3, which gives Method 4 a score of 0 and Method 3 a score of -1.

With connecting flights, Methods 3 and 4 would provide easy ways to integrate connecting flights without significant additional work, earning a score of +1. Method 2 earns a -1 because it would require developing a way to define the maximum EEF for a route.

Integration of non-contrail SLCPs would be relatively easy for all methods which gives them all +1 scores. Understandability of the metrics is an open question that would benefit from further debate at the AC level.

Summing up the scores, Methods 3 and 4 would be in the lead with scores of 3 and 4. Method 2 is lagging with a score of 0.

FUTURE WORK

Beyond choosing a method for communicating the impact, there is further work for the Task Group. Specifically:

- Revisit the threshold definitions within the methods to assess if there is a better way to define the buckets.
- Integrate the impact of non-contrail SLCPs into the communication of warming impact.
- Extract the dependence of contrail impact for day versus night flights from this analysis.
- Produce a guidance document on how to use the TIM output. While it is not possible to control how the TIM output is used by booking platforms, a guidance document could help communicate how the output should be used.

Beyond the specific work of this workstream, here is a list of issues that we would like to address but will not be able to in the time frame of this workstream, because there is not yet sufficient data:

- **Validating the climatological model against meteorological model results:** It is essential to validate the predictions made by this analysis against a meteorological analysis that uses reanalysis data, which represents our best estimate of the real-world weather conditions. This would help to understand if making classifications based on a climatological model is an appropriate way to communicate the contrail climate impact of specific flights.
- **Receiving customer feedback:** We would like to better understand how our efforts at communicating the climate impact of flights are being received by consumers. This can only happen after the new impact estimates are released and the public is able to see them and use them.
- **Dependence of contrail impact on longitude:** We are capturing trends based on flight latitude but not longitude. It is known that surface albedo plays a role in the radiative balance of incoming and outgoing radiation. Incorporating the impact of longitude would be needed to capture these effects. However, this would require further research.
- **Dependence of contrail impact on aircraft type:** Research shows that contrail impact can change based on the aircraft and engine type. This nuance is not being captured in the current analysis, and more data is required to reliably extract trends in contrail impact based on aircraft type.

APPENDIX

Here we present a few additional analyses that are supporting investigations and are not central to the choice of the TIM output. As such, this section is less refined.

Bucket definition thresholds

The table below presents a few different options for bucket definitions, along with the distribution of the flights in each bucket.

Table A1. Options for bucket definitions for the CO_2 multiplier (Method 4)

Description	Options for buckets	Percentage of flights
Original buckets	• Cooling ($x_i < 0$) • Negligible ($0 \le x_i < 0.2$) • Warming ($0.2 \le x_i < 1.2$) • Strongly warming ($1.2 \le x_i$)	• Cooling: 0.27% • Negligible: 33% • Warming: 47% • Strongly warming 20%
Spreading the negligible bucket over -0.1 to 0.1	• Cooling ($x_i < -0.1$) • Negligible ($-0.1 \le x_i < 0.1$) • Warming ($0.1 \le x_i < 1.$) • Strongly warming ($1. \le x_i$)	• Cooling: 0.007% • Negligible: 22% • Warming: 52% • Strongly warming 26%
Creating an equal distribution across bins	• Cooling ($x_i < -0.2$) • Negligible ($-0.2 \le x_i < 0.2$) • Warming ($0.2 \le x_i < 0.7$) • Strongly warming ($0.7. \le x_i$)	• Cooling: 0.001% • Negligible: 33% • Warming: 34% • Strongly warming 33%

Cooling flights

There is a negative contrail impact on 0.26% of flights, but some of these are very close to 0.

We should extend the negligible impact bucket to negative values, as well. Here are the statistics with different limits for the cooling flights.

Table A2. Options for the cooling bucket definition for the CO_2 multiplier

Definition of cooling bucket	Percentage of flights	Number of origin-destination pairs represented	Seasons
< 0 contrail impact	0.27%	1,060	All
< -0.1x contrail impact	0.009%	71	Only summer and autumn
< -0.2x contrail impact	0.001%	15	Only summer and autumn

Below is the plot of all the routes that have $x_i < -0.2$

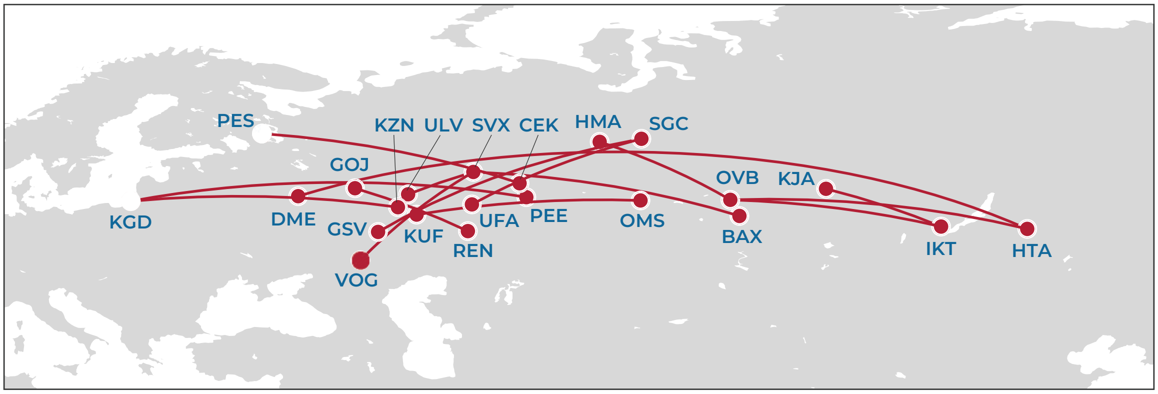

Figure A1. All routes that have CO_2 multiplier $x_i < -0.2$

Seasonal differences

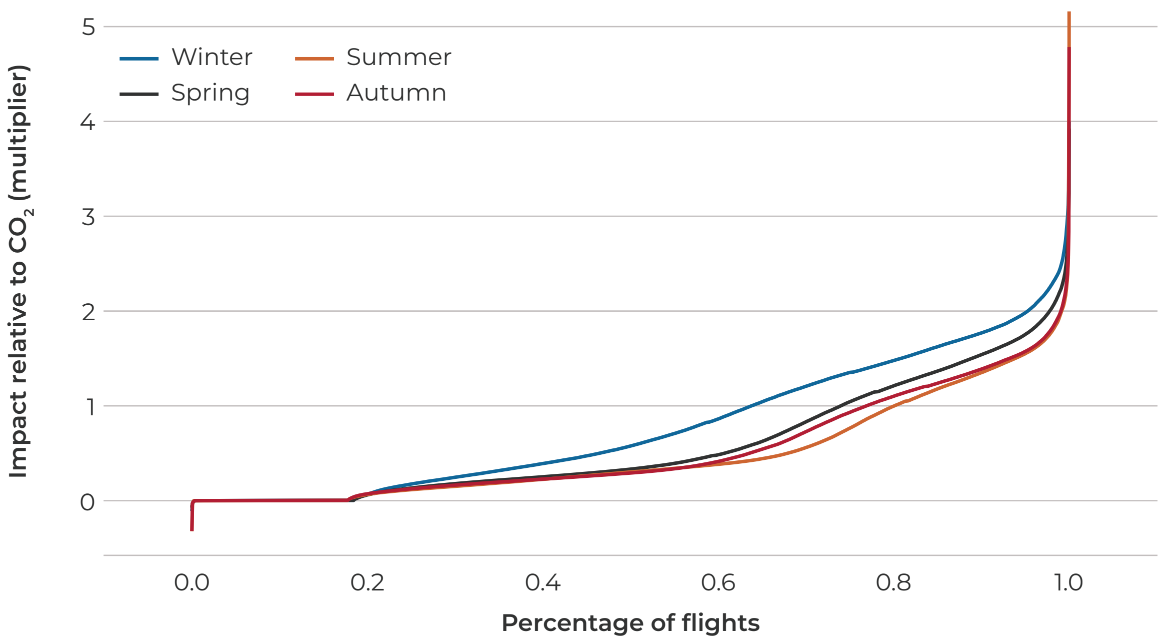

Figure A2. Seasonal variation in the distribution of the CO_2 multiplier required to represent contrail climate impact. Flights were first sorted in ascending order of the multiplier and then plotted on this figure.

There is a clear difference in the warming impact with the season. Winter generally produces the most warming impact, followed by spring and then autumn. The summer season generally produces the least warming impact.

Table A3. Statistical metrics for the seasonal distribution of the CO_2 multiplier

	Median	Minimum	Maximum
Winter	0.58	-0.09	3.98
Spring	0.33	-0.06	3.90
Summer	0.30	-0.30	5.14
Autumn	0.29	-0.30	4.76