F
Lying is
A
Privilege

# Modelling a Non-Judgemental Approach to Air Travel Reduction

Richard Morrey
23rd June 2020
https://flapyourarms.org

# Contents

# 1 Introduction

## 1.1 Goal

Until a commercially viable, carbon-neutral propulsion system for aircraft is developed and widely adopted there is only one way to tackle the increasing contribution of air travel to Global Warming: the current exponential growth in passenger numbers[1] must come to an end[2]. This paper explores how a specific non-judgemental scheme might achieve such a goal; summarising the approach, it's implementation, and the outcome of a computer modelling exercise to evaluate its viability.

## 1.2 A Non-Judgemental Approach

An annual holiday in the sun has become a central part of the developed lifestyle. Foreign travel is actively encouraged as essential to personal fulfillment, and the more frequent and exotic it is the better. The freedom to travel easily to any location in the world is taken for granted by many and seen as a right by some.

[1] World Bank, 2020. "Air Transport, passengers carried". https://data.worldbank.org/indicator/IS.AIR.PSGR. Last accessed 20th June 2020

[2] Lucas, Caroline.2019. "Heathrow's expansion plans make a mockery of the zero emissions strategy". The Guardian. 18th June 2019. https://www.theguardian.com/commentisfree/2019/jun/18/heathrow-expansion-mockery-zero-emissions . Last accessed 21st July 2019.

In democratic and free societies policy makers look to the tax system as the primary means of changing the behaviour of their citizens, but new taxes are always controversial. The "gilets jeune" protests in France show what ill-judged attempts to levy a carbon tax on fuel can lead to[3]. In the UK a "frequent flyer levy" has been proposed, in which individuals pay more tax per flight the more often they fly[4]. Whilst a progressive tax like this is less likely to be perceived as unfair it still carries the risk of being counterproductive through stigmatisation.

Stigmatising a behaviour can lead to resentment and entrenchment rather than positive engagement with behaviour change. In the case of a frequent flyer levy this may drive actors to simply pay the tax and carry on flying. Non-judgemental approaches to behaviour change originated in the treatment of drug addicts[5]. They start with the principle that the actor is not to blame, and that any approach should be consistent with this absence of blame. Drug addiction is approached as an affliction to be treated rather than a misbehaviour to be punished.

Psychological studies have shown that even positive financial incentives can be counterproductive, by eroding the moral imperative[6] to change behaviour. The approach to air travel reduction presented here, Flying is a Privilege (FLAP), not only aims to be non-judgemental but is also free of any kind of financial lever, positive or negative.

[3] Wilsher, Kim. 2018. "Macron scraps fuel tax rise in face of gilets jaunes protests". https://www.theguardian.com/world/2018/dec/05/france-wealth-tax-changes-gilets-jaunes-protests-president-macron . Last accessed 28th July 2019.

[4] A Free Ride. http://afreeride.org/. Last accessed April 2020.

[5] Rafaeli, J.S. with Woods, Neil. 2019. "Fighting the Wrong War". This is not a Drill: An Extinction Rebellion Handbook. Penguin Random House UK.

[6] Raworth, Kate.2017. Doughnut Economics. Penguin Random House.pp118-120.

## 1.3 How it works

The FLAP scheme works as follows:

- A distance account, analogous to a bank account, is opened for every traveller. Each account has a balance in kilometres, which starts at zero.
- Airlines withdraw the distance of each flight from the traveller's account at check-in.
- The "daily allowance" defines a total distance in kilometres. The account of every overdrawn traveller is credited with an equal share of the Daily Allowance at the start of each calendar day.
- Overdrawn travellers are not allowed to start a new trip, with airlines refusing any attempt to check in.

Lowering the daily allowance over time has the effect of lengthening the time between trips and thereby reducing air travel. Sharing the daily allowance out between those travellers in deficit is equitable but at the same time allows for continuing variability in the distance flown, both by individuals and overall over the course of each calendar year.

The concepts of “trip” and “clearance promise” add the additional flexibility needed to make the scheme practicable and low impact for all types of traveller. They work as follows:

- A trip is defined as any sequence of consecutive flights by a single traveller with a maximum length of 90 days from the start of the first flight to the end of the last.
- Travellers (or the airlines they fly with) are able to submit the details of a planned trip to FLAP in advance and obtain a clearance promise. This is a committed date after the completion of a specific trip when the traveller is allowed to fly again. regardless of their distance balance.
- Consecutive clearance promises can be "chained" to allow clusters of trips close together in time. Chained promises have a date that is artificially brought forward to accommodate the next trip. Chains have a maximum length, with the date of the last promise pushed further into the future to compensate for the distance deficit accrued over the length of the whole chain.

The overall effect of the system is to enforce a decreasing limit on how much distance is flown:

- Without placing a direct restriction on how much any one person can fly.
- Whilst applying the same rules apply to all travellers, regardless of how often they fly.
- Without levying taxes or making use of any other kind of financial incentive.

## 1.4 Usability

Clearance promises remove the need for travellers to understand how the scheme works. Instead they interact with a simple booking calendar as illustrated below for a maximum promise chain length of three. In **Figure 1** the traveller has booked a first trip in February. The trip is green to indicate there is plenty of scope for booking further trips.

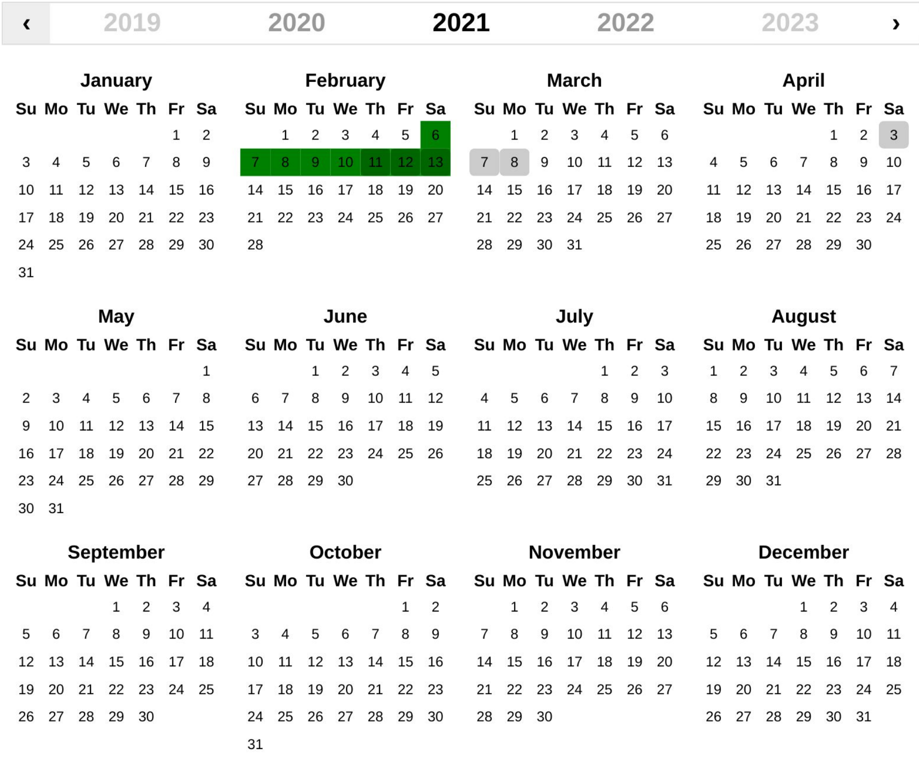


**Figure 1 Planning the first trip**

In **Figure 2** the traveller has booked a second trip in April. Both trips are now orange to indicate the calendar is starting to fill up. However all other dates are still available.

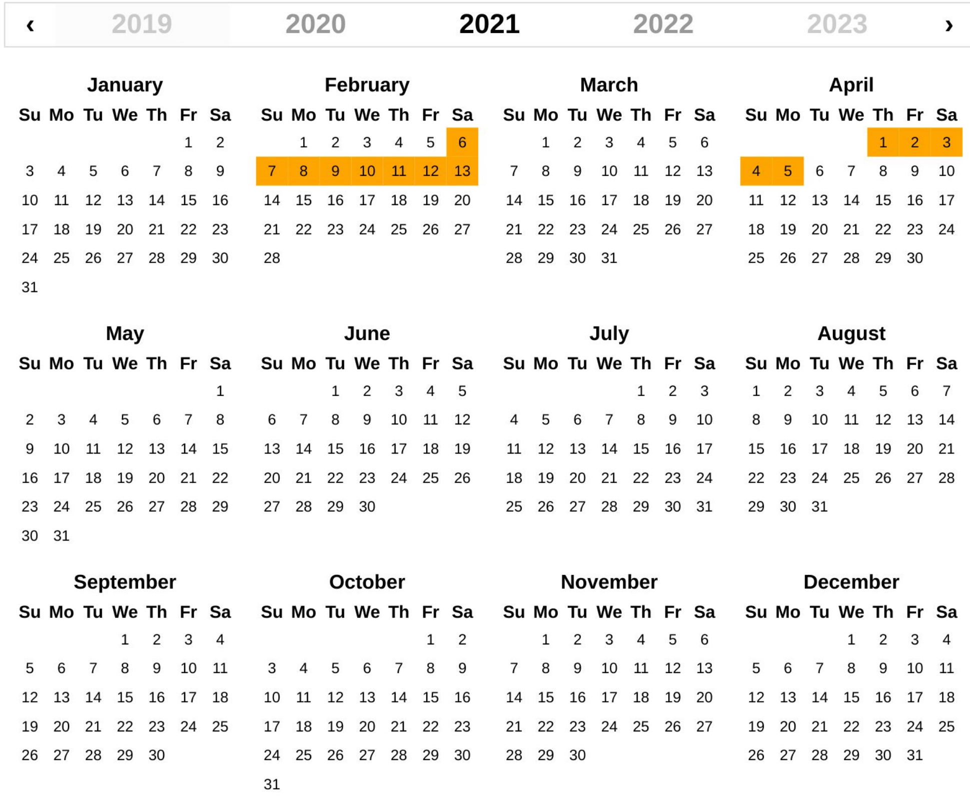


**Figure 2 Planning the second trip**

In **Figure 3** the traveller has booked a third trip in June. The promise chain has now reached its maximum length so all trips are colored red indicating no more trips can be added until after the clearance date of the last trip in the chain is set. All unavailable dates are greyed out.

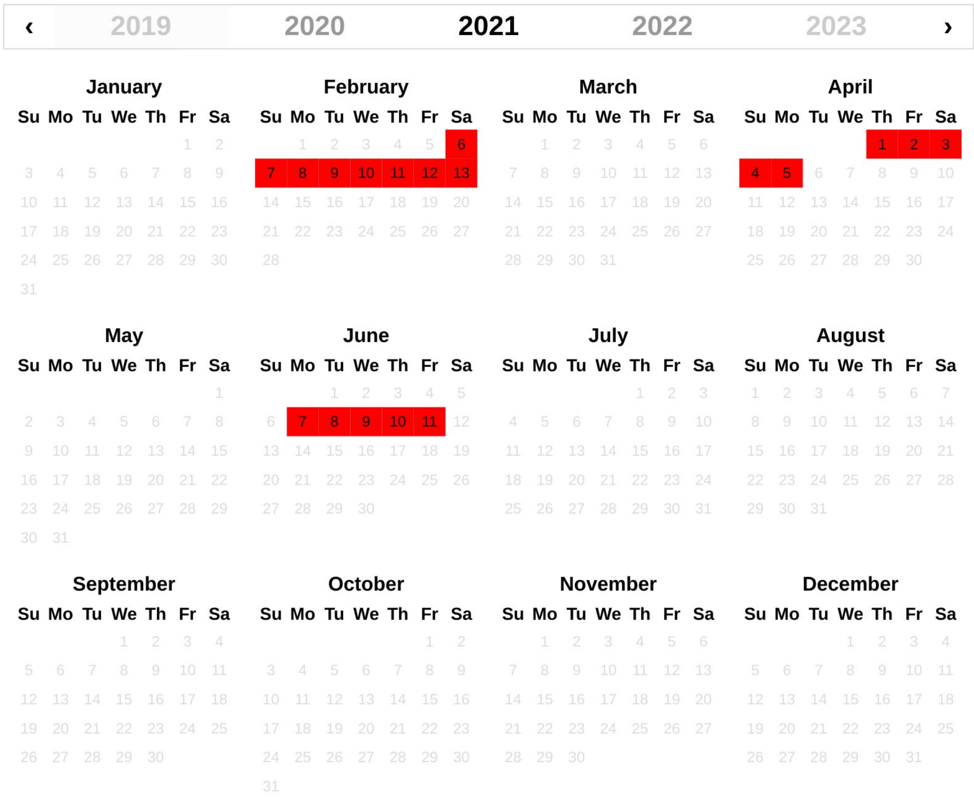


**Figure 3 Planning the third trip**

There are other scenarios as trips are taken and promise dates calculated, but these can also be presented to the user as a simple calendar showing available dates for travel.

# 3 Computer Modelling

## 3.1 Aims and Approach

Computer modelling was carried out with the aims of:

1. Establishing the viability of implementing and running FLAP as a global scheme.
2. Demonstrating the effectiveness of FLAP in reducing air travel
3. Investigating the impact of FLAP on the traveller experience.

To achieve these aims a full and independent FLAP system was implemented, exercised by an external modelling tool. Also the focus was on modelling those aspects of air travel most likely to negatively impact effectiveness and the traveller experience. No attempt was made to generate anything other than a basic approximation of real world global traffic flows.

The source code for both FLAP and the modelling tool (with usage instructions) are publically available[7].

## 3.2 FLAP System

The FLAP system is a Go[8] package that exposes the following public methods:

[7] See https://github.com/richardmorrey/flap. Last accessed

[8] Go. https://golang.org/. Last accessed 28th April 2020.

- `SubmitFlights`. Used by airlines to report the check-in by a specified traveller for a specified sequence of one or more flights. Travellers are identified by passport number and country of origin code. Flights are defined as two ICAO airport codes[9] and a start and end date/time.
- `UpdateTripsAndBackfill`. Invoked once a day to update the status of all travellers: sharing out the dally allowance, "keeping" clearance promises, and enforcing the ending of trips that have exceeded 90 days in length.
- `Propose` and `Make`. Used by travellers to obtain clearance promises. The traveller invokes `Propose` to obtain a clearance promise proposal for a specified trip. The traveller can choose to accept the proposal by invoking `Make`, at which point FLAP is committed to keeping the promise at the point when the specified trip ends.

`SubmitFlights` returns an error if the specified traveller is overdrawn or does not have a kept clearance promise - an instruction to the submitting airline to disallow the check-in.

A database table of longitude and latitude coordinates for all ICAO airport codes allows the physical distance between the start and end airports of each submitted flight to be estimated using the haversine formula[10]. A further fixed distance is

[9] Assigned by the International Civic Air Organization and published as document ICAO 7910.

[10] https://www.geeksforgeeks.org/haversine-formula-to-find-distance-between-two-points-on-a-sphere/. Last accessed 4th May 2020.

added to account for the fuel used during taxiing and take-off[11] to give a final value to use in all subsequent calculations.

For clearance promises to be effective they need to provide consistent and accurate predictions of how long it will take for a traveller's balance to return to credit following the end of a trip. The approach taken here is to perform linear regression against daily allowance share over time. Three inputs into the regression calculation are configurable: time frame (in days), rolling average window (in days), and the degree of the resulting polynomial.The most important variable for clearance promises, the number of days in advance of the end of a trip a promise can be given, is also configurable.

## 3.3 Modelling Tool

The modelling tool is also implemented as a Go package. It takes publically available lists of global flight paths and airports[12], and constructs from that data:

- A set of countries, each holding a set of airports and weighted according to the size of its airport set.
- Sets of airports, each with a set of flight paths and weighted according to the size of its flight path set.
- Sets of flight paths, each containing a destination airport and weighted according to the weight of the destination airport.

[11] Jung, Yoon. 2010. "Fuel Consumption and Emissions from Airport Taxi Operations". Green Aviation Summit NASA Ames Research Center, Sept 8-9, 2010. https://flight.nasa.gov/pdf/18_jung_green_aviation_summit.pdf . Last accessed 21st July 2019.

[12] https://openflights.org/.

Simulated travellers are assigned a country using weighted probabilities against the set of countries. Trips are chosen using weighted probabilities to select a departure airport and then flight path from the traveller's home country. All trips are simple return flights with no onward connections.

Thus while no attempt is made to accurately simulate real world flight patterns, the model does follow real-world flight paths, with flights chosen to give an impression of real-world traffic flow.

Traveller behaviour is modelled using a series of configurable bands, with each band defining a proportion of the total traveller population with the same probability of planning travel on any given day. For each day of the model planning weight probabilities are used to decide which travellers are planning travel. Those that are then:

1. Choose a trip length, departure airport, and destination airport, all using weighted probabilities.
2. Randomly choose a future start date that fits within the allowed range for clearance promises.
3. Request a clearance promise for the planned trip. If this can't be provided given the traveller's current schedule the trip is recorded as "cancelled". Otherwise the promise is made and the trip planned in.

Real world traffic flows vary seasonally. The model simulates this by allowing the allocation, to each traveller band, a set of twelve monthly weights. These are applied such that the probability of starting a trip each day varies according to the calendar month but the mean of which over the course of each calendar year is equal to the specified planning probability of the band.

The model runs for a configured period of days. For each of those days, for the entire traveller population, it: invokes `Propose` and `Make` to plan Trips; invokes `SubmitFlights` for due flights; and invokes `UpdateTripsAndBackfill` to enforce trip completions, apply clearance promises, and share out the daily allowance.

## 3.4 Configuration

The outcome presented here is for a single execution of the model aiming to emulate real world traffic flows. With little information in the public domain about real world traveller behaviour, some assumptions were needed to come up with an appropriate configuration.

A 2014 UK Department for Transport passenger survey[13] includes a breakdown of flights taken per passenger over the preceding 12 months, see **Table 1**. This useful but incomplete information was adapted to calculate planning probabilities for six traveller bands[14], as shown in **Table 2**. These bands were distributed over a total traveller population of 1 million.

[13] DFT. 2014. Public experiences of and attitudes towards air travel: 2014. https://www.gov.uk/government/statistics/public-experiences-of-and-attitudes-towards-air-travel-2014. Last accessed 26th May 2020.

[14] The midpoint of each Trip Frequency range was used as input to the function.

| **Return flights in last 12 months** | 0 | 1 | 2 | 3 | 4 or more |
|---|---|---|---|---|---|
| **% of survey respondents** | 51 | 22 | 11 | 5 | 10 |

**Table 1 UK DFT Passenger Survey Results 2014** [15]

| **Band** | **Trips per annum** | **% Traveller population** | **Trips per day** |
|---|---|---|---|
| 1 | 0-1 | 50 | 0.0014 |
| 2 | 1-2 | 25 | 0.0041 |
| 3 | 2-3 | 12 | 0.0068 |
| 4 | 3-4 | 6 | 0.0096 |
| 5 | 4-6 | 5 | 0.014 |
| 6 | 6-12 | 2 | 0.018 |

**Table 2 Traveller band configuration**

[15] *Ibid,* Figure 1.

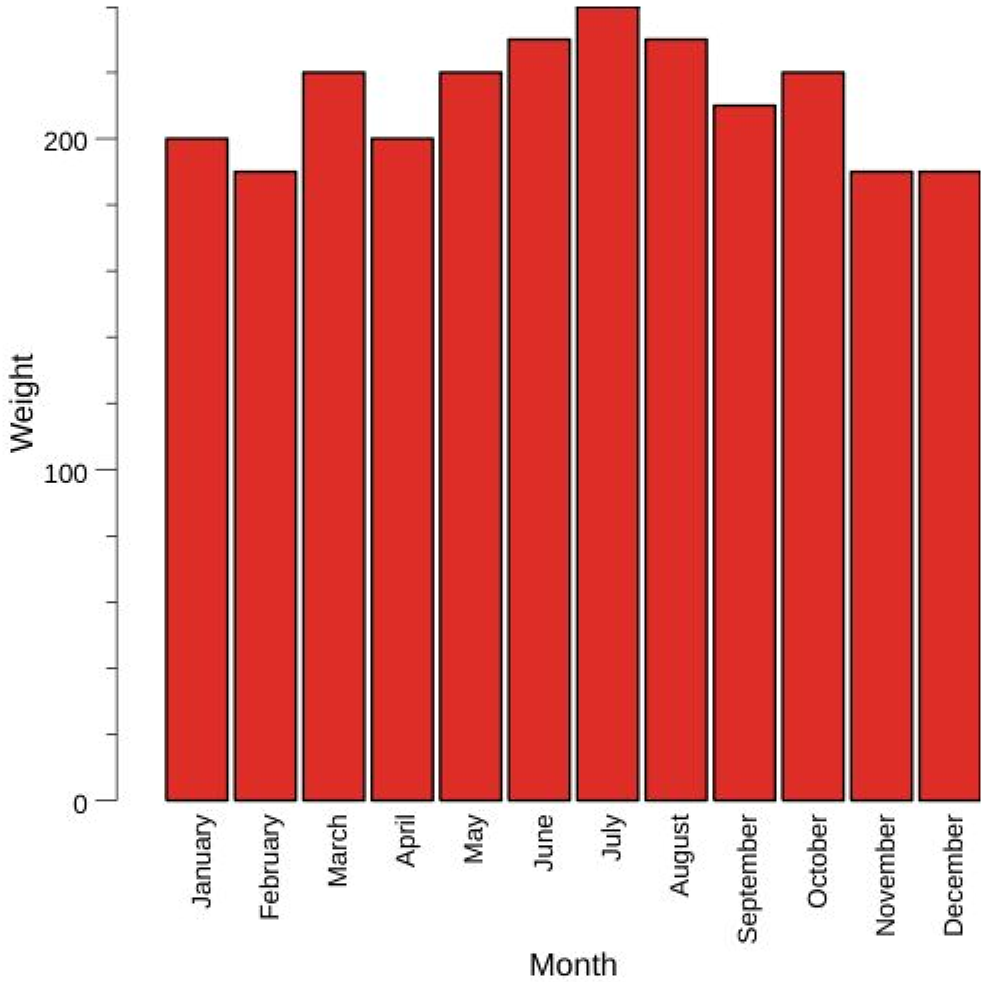


**Figure 4 Monthly planning weights configuration**[16]

[16] Liang Mao, Xiao Wu, Zhuojie Huang, Andrew J.Tatem.'Modeling monthly flows of global air travel passengers: An open-access data resource'. https://www.sciencedirect.com/science/article/pii/S0966692315001581. Figure 4b. Last accessed 12th May 2020.

Each band had the same set of monthly weights applied to create a realistic variation in flight volumes over the course of each calendar year. In the absence of any real world data, these were derived from a detailed research model of global traffic flows, as shown in **Figure 4**.

Clearance promises were configured such that promises were provided up to 90 days in advance, and using a simple linear regression against data for the preceding 100 days in order to predict dates.

The model ran for a total of 2000 days. For the first 365 of these a daily allowance was not enforced. On Day 366 the mean distance travelled over the preceding days was used as the initial daily allowance. Thereafter the daily allowance was reduced by 0.01% each day.

# 4 Results

## 4.1 Viability

The model ran on an 8 core Intel virtual machine with 64 Gb of physical memory and 100 Gb of SSD disk storage. It took 45 hours to complete at a cost of USD 18.

The model simulated approximately 0.1% of global travellers[17], and 2000 days of travel. Thus, assuming linear scaling, it simulated the equivalent of one day's global travel every 23 hours, meaning the monthly compute cost for all global travellers would be of the order of USD 1000.

The total size of the traveller database at the end of the simulation was 750 Mb, meaning a database for all global travellers would be of the order of 1 Tb with a monthly cost of the order of USD 1000[18].

These numbers are only indicative, and there are other concerns, discussed in a previous white paper[19]. However from a scalability and cost perspective a global FLAP system is certainly viable.

[17] Assuming one billion travellers in total. There is no clear information in the public domain on what this number actually is.
[18] Based on estimate from Google Cloud Platform Pricing Calculator using Bigtable. https://cloud.google.com/products/calculator. Last accessed 20th June 2020.
[19] Morrey, Richard. 'FLAP: A Simple, Fair and Viable Scheme for Reducing Air Travel' http://www.flapyourarms.org/doc/flapwhitepaper.pdf. Ch 4. Last accessed 20th May 2020.

## 4.2 Effectiveness

**Figure 5** shows how total distance travelled by the traveller population (red line) varied over time as the daily allowance was reduced (black dashed line). It shows a clear and consistent reduction over time such that air travel at the end of the 2000 day period had been reduced by 40%.

The repeating "shadow" of **Figure 4** in the red line confirms that this overall reduction was achieved whilst continuing to allow significant variation in traffic volumes during each calendar year.

Total distance travelled consistently exceeds the daily allowance over time. This is due to the accumulation of deficits by travellers who are mid-trip or who are part way through a clearance promise chain. It is the concept of distance overdrafts that gives FLAP the flexibility to support variable traffic flows using a very simple algorithm.

## 4.3 Traveller Experience

**Figure 6** shows the average percentage of trips cancelled per traveller over time, broken down by band as defined in **Table 2**. The darkness of the shade varies according to the frequency of travel - so Band 1 is white and Band 6 is the darkest shade of red.

The percentage of cancellations increases both with time, as the distance allowance is reduced, and band, as the planned trip frequency increases. Band 1 and Band 2, representing 75% of the traveller population, suffer very little disruption to their plans. In the case of Band 1, no trips are cancelled until Day 700 of the modelling period and even at the end less than 2% are being cancelled.

**Figure 7** shows how the total distance travelled per traveller for each of the six bands varied. It shows that, despite an inevitable narrowing over time, variation in distance travelled between bands is maintained throughout the modelling period. Even on Day 2000 travellers in Band 6 still have the freedom to travel three times further per day than those in Band 1.

**Figure 8** shows how closely the linear regression used for calculating clearance promises (blacked dash line) tracked distance allowance share for Day 1000 and Day 2000. The regression achieves a close match, and this is typical of all days.

**Figure 9** shows the distribution of traveller balances at the point when the clearance promise date was reached[20]. A perfect clearance promises algorithm would have a very tight distribution of values around zero. Whilst share decay is approximately linear over short intervals it decays exponentially over the entire modelling period so inaccurate do occur, particularly for long trips with old promises. These inaccuracies cause distance deficits for frequent travellers to accumulate over time.

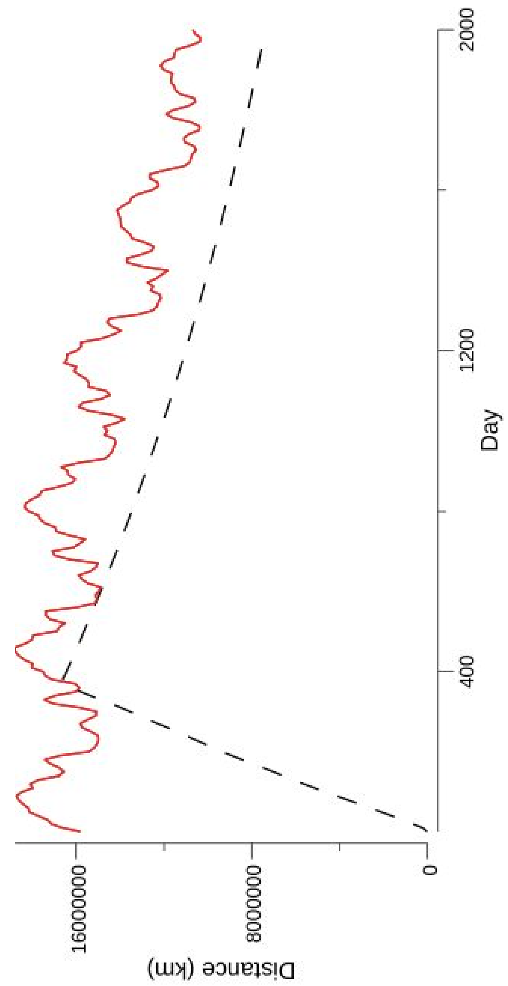


**Figure 5 Total distance travelled and daily allowance over time**

[20] Note the distribution is heavily truncated on the positive side because Travellers who are in credit do not receive a share of the Daily Allowance.

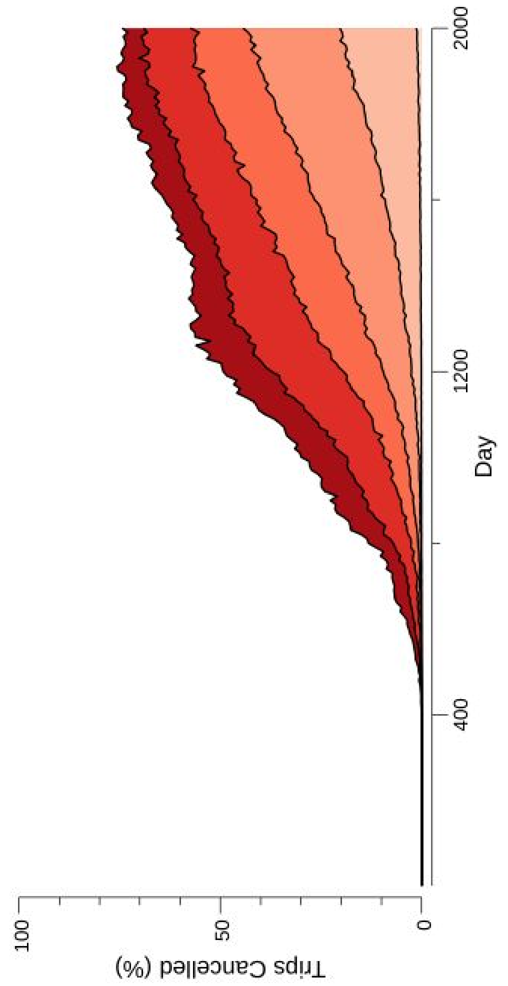


**Figure 6 Percentage of trips cancelled per traveller over time**

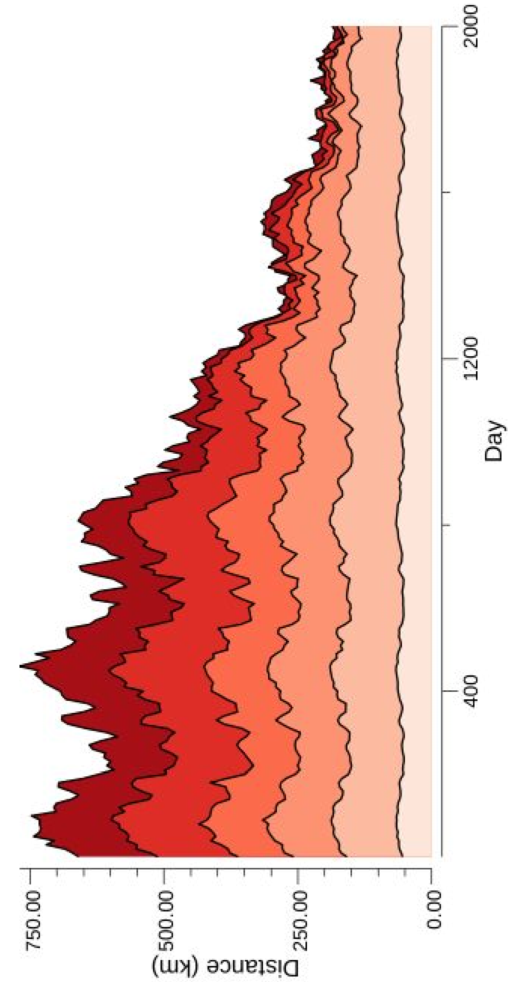


**Figure 7 Mean distance per traveller per day by band**

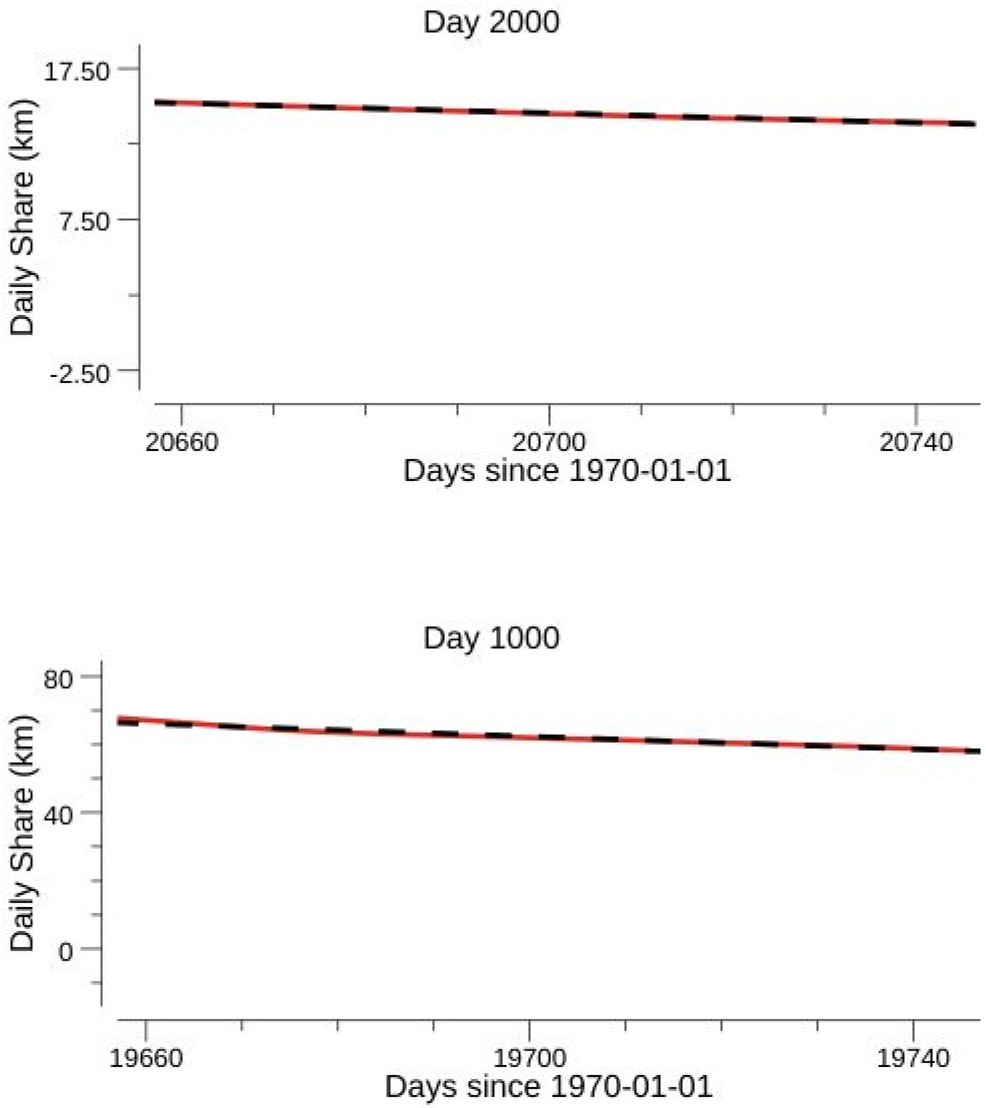


Figure 8 Linear regression against daily share

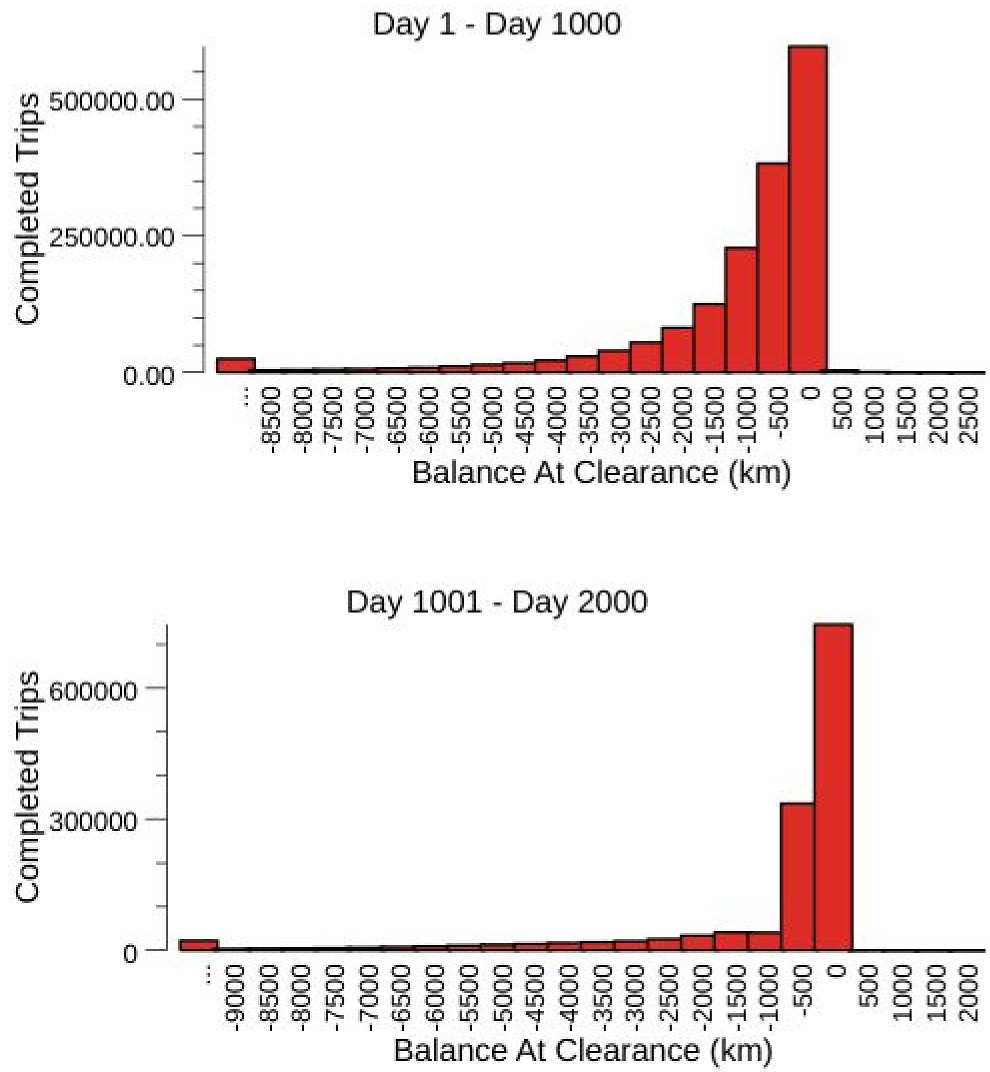


Figure 9 Distribution of traveller balance at clearance

# 5 Conclusions

The scenario modelled here is an extreme one in that the daily allowance is reduced on a daily basis for five years, with traveller behaviour remaining unchanged throughout – a stress test. Even in this case the results show the system working equitably and effectively, reducing air travel in-line with the daily allowance whilst continuing to allow variability in seasonal travel flows and individual traveller behaviour.

The imperfections of the clearance promises algorithm are clearly an area for improvement. The leveraging of more advanced machine learning approaches would likely result in more responsive and accurate predictions. However the results show that even with an imperfect solution the overall outcome is good.

The only distance allowance strategy considered here is one of gradual reduction. However in a real world implementation the strategy could be tailored to match political realities, effecting a different rate of decrease, a steady state, or even merely to moderate increase in air travel over time. This is a level of control that not even a progressive and carefully calibrated tax-based approach could provide.