D5.2 Final Assessment Report

Deliverable 5.2

Vista

Grant:	699390
Call:	H2020-SESAR-2015-1
Topic:	SESAR-05-2015 Economics and Legal Change in ATM
Consortium coordinator:	University of Westminster
Edition date:	09 November 2018
Edition:	01.02.00

Authoring & Approval

Authors of the document

Name/Beneficiary	Position/Title	Date
Luis Delgado / University of Westminster	Consortium Member	05 November 2018
Gérald Gurtner / University of Westminster	Consortium Member	05 November 2018
Samuel Cristóbal / Innaxis	Consortium Member	05 November 2018
Jorge Martín / Innaxis	Consortium Member	05 November 2018
Andrew Cook / University of Westminster	Project Leader	05 November 2018

Reviewers internal to the project

Name/Beneficiary	Position/Title	Date
Graham Tanner / University of Westminster	Consortium Member	09 November 2018

Approved for submission to the SJU By — Representatives of beneficiaries involved in the project

Name/Beneficiary	Position/Title	Date
Andrew Cook / University of Westminster	Project Leader	09 November 2018

Rejected By - Representatives of beneficiaries involved in the project

Name/Beneficiary	Position/Title	Date
N/A		

Document History

Edition	Date	Status	Author	Justification
01.00.00	16 July 2018	Release	Vista Consortium	New document for review by the SJU
01.01.00	14 September 2018	Release	Vista Consortium	Incorporating SJU comments
01.02.00	09 November 2018	Release	Vista Consortium	Incorporated SJU comments

Vista

MARKET FORCES TRADE-OFFS IMPACTING EUROPEAN ATM PERFORMANCE

This deliverable is part of a project that has received funding from the SESAR Joint Undertaking under grant agreement No 699390 under European Union's Horizon 2020 research and innovation programme.

Abstract

Vista examines the effects of conflicting market forces on European performance in ATM, through the evaluation of impact metrics on four key stakeholders, and the environment. Vista models the current, 2035 and 2050 timeframes based on various factors and their potential evolution. Vista's model covers the three temporal phases of ATM (strategic, pre-tactical and tactical), and represents a typical (busy) day of operations. The model is able to estimate the impact of factors on the different phases independently, allowing us to capture how indicators change under different scenarios and execution phases. This deliverable presents the final results obtained from the model, together with a detailed description of the various parts of the model, the analyses performed to prepare the data, and the model calibration.

Table of Contents

Table of Tables

Founding Members

Table of Figures

Executive summary

Vista examines the effects of conflicting market forces on European performance in ATM, through the evaluation of impact metrics on four key stakeholders, and the environment. Vista models the **current, 2035 and 2050 timeframes** based on various factors and their potential evolution. Vista's model covers the three temporal phases of ATM (strategic, pre-tactical and tactical), and represents a typical, busy day of operations. The model is able to estimate the impact of factors on the different phases independently, allowing us to capture **how indicators change under different scenarios and execution phases**.

This deliverable presents the final results obtained from the model, together with a detailed description of the various parts of the model, the analyses performed to prepare the data, and the model calibration.

Section 2 describes the developed model in detail. The approaches selected for the various layers in the model are presented. The strategic layer implements an **agent-based economic model** and a **schedule mapper** to generate demand and capacity for the different stakeholders, flight schedules and passenger flows – it feeds the pre-tactical layer. The pre-tactical layer transforms the output of the strategic layer into individual flight plans, **passenger itineraries and ATFM regulations** delay – it feeds the tactical layer. The tactical layer runs the day of operations, tracking flights and passengers and reacting to the tactical situation in the system.

The model includes estimates of **door-to-door times**, for example showing the impact of flight delays on door-to-door times, which is an important capability set in the context of assessing the Flightpath 2050 vision of 90% of EU passengers travelling within four hours door-to-door for journeys involving air as one of the modes.

The **cost of uncertainty** is also estimated, in euros, by calculating the difference of the cost of delay in two situations, and included as a new assessment metric under the various scenarios. This is a valuable feature of the model, since SESAR envisions substantial reductions of uncertainty in the future, the impacts of which Vista can quantify economically.

In Section 3, the calibration activities that have been performed are presented. A significant amount of effort has been devoted to this task to ensure that the model is adequate and captures the impact of factors properly.

The scenarios and factors modelled are described in Section 4. Vista deploys five background scenarios: 'current', 'L35' (low economic growth and low technological development in 2035), 'H35' (high economic growth and high technological development in 2035), and likewise 'L50' and 'H50' for 2050. The scenarios also include 'supportive' and 'non-supportive' contexts, modelled through five foreground factors: the **price of fuel**; implementation level of **4D trajectories**; level **of passenger reaccommodation** tools; level of **(regulatory) passenger** provision schemes; and, level of **(regulatory) emissions** scheme.

The background scenarios L35 and L50 are considered particularly relevant and are the ones for which different foreground factors will be analysed. Baseline H35 and H50 are modelled to push the system out of its current operational environment. A total of nine scenarios have been modelled across the different layers of the model:

- five baselines:
  - current, low 2035, high 2035, low 2050, high 2050;
- and four scenarios built on top of the low 2035 and low 2050 baselines:
  - non-supportive 2035, supportive 2035, non-supportive 2050 and supportive 2050.

The final results are described layer-by-layer in Section 5, with detailed analyses. Key metrics and results are presented at the end of each layer's section. Key metrics investigated also include **the cost of delay**. Considering its importance in airline operations, it is surprisingly rare to have an estimation of this cost taking the full distribution of delays into account. The **total delay per flight** shows how much the system will be under stress in the future; if the airports are indeed the main contributors to delay, we show that the ANSPs, in turn, have a high potential for greater delays, since they move close to their maximum capacities. The **total level of emissions in CO_2 equivalents** has an obvious importance, given rapid climate change; the emissions are all too often only measured per flight, which does not show their full impact. The model predicts a **substantial increase in the total emissions** due to the increase in traffic, differences in the emissions per flight expected can be observed across the different layers of the model due to the different assumptions made strategically, pre-tactically and the actual fuel consumptions computed tactically. The key results per layer are indicated below.

Key strategic results.

1. The main drivers for most of the metrics are the demand and the price of fuel. See deliverable D5.1.
2. ANSPs may get close to their maximum capacity (set to 120% of current capacity, with technological advancements), depending on the scenarios, and trigger some significant delay for airlines, even with the technological advancements envisioned by SESAR. See Sections 5.1.2.2 and 5.1.3.2.
3. ANSPs see their **unit rates decrease substantially** in the future. This is due to the joint effect of higher levels of traffic and greatly improved efficiencies. See Section 5.1.2.2.
4. Airports create most of the delay, and the **increase of capacity envisioned by SESAR is not sufficient** on its own to deal with the increase in traffic. See Section 5.1.3.1.
5. The cost of emissions only really has an impact on airlines when **NO_x is taken into account** together with a great increase in the price of allowances. Emissions have otherwise (almost) no impact on the cost structure of the airlines. See Section 5.1.3.1.
6. The average size of the aircraft used by airlines is increasing. This has an impact on the average cost of fuel and environmental impact per flight, despite various other improvements such as the length of flights. Other technological improvements related, for instance, to engine efficiency have not been considered in Vista. See Section 5.1.3.1.

7. Total emissions are expected to very substantially increase in the future. They are mainly driven by the increase in traffic, and, to a lesser extent, by the increase in the average size of the aircraft. See Section 5.1.3.1.
8. The reduction of uncertainty on the departure time envisioned by SESAR is expected to a have major impact on the cost of delay to the airlines. The **cost of uncertainty represents roughly half of the total cost of delay in the current scenario**, but its share drops substantially in the future. See Section 5.1.4.1.
9. Passengers usually see a moderate decrease of fare with respect to their income, except when the operational cost of the airlines increases too much, notably because of the price of fuel. See Section 5.1.3.2.

Key pre-tactical results (see Section 5.2.3).

1. Fuel consumption per flight is flat over time as the (e.g. technological) benefits obtained by the system are offset by the **use of longer routes with larger aircraft**, with a potential shift to greater fuel efficiencies. The relative importance of the fuel price over time might also favour the selection of trajectories that use less fuel.
2. The selection of larger aircraft over time is related to the increase of passenger demand and route length.
3. There is an **increase in the size of the buffers** per flight: this may contribute to the **reductions in tactical delay costs** and could be used by the strategic layer to tighten the schedules – the reason for this increment may be linked to having more numerous longer routes, which usually have larger buffers to manage greater uncertainty.
4. The number of passengers connecting increases over time – these effects are also discussed further in the tactical results.

Key tactical results (see Sections 5.3.1 and 5.3.2).

1. Most passenger- and flight-centric metrics follow similar trends overall, but there are non-linear differences on how metrics scale, i.e., **no simple, direct translation between passenger and flight metrics**.
2. Reducing delay, either departure or arrival, has a limited effect on the total door-to-door travel times for passengers.
3. Reductions in flight arrival delay with passenger arrival delay map close to a 1 : 1.3 ratio. That is, on average, one minute of flight delay corresponds to 1.3 minutes of delay per passenger. This is due to the fact that the delay experienced by passengers is higher due to missed connections.
4. There is a **diminishing return of the positive effects** (in terms of delay and environmental impact) **of foreground factors** (implementation of 4D trajectories, advanced passenger reaccommodation tools, protective regulatory passenger provision schemes and regulatory emissions schemes) in the long term (e.g. 2050).

5. There is a clear **trade-off between delay performance and cost metrics**: improving system performance is usually expensive; Vista can quantify such trade-offs to find compromise solutions.
6. The results show that an improvement in passenger door-to-door times does not necessarily imply an increase in the average emissions per flight.
7. The average emissions per flight tend to slightly decrease, remaining mostly flat, but the increase in traffic leads to an overall higher impact of aviation on the environment.

The modularity of the development of the model layers allows the enhancement or part replacement of the model seamlessly. Vista is capable of capturing and quantifying the relationship between complex metrics across several stakeholders, reproducing classic KPIs and estimating complex and newly defined metrics.

Vista is unique in the sense that it supports the analysis of how a given metric changes during the temporal evolution of the different ATM phases: from the expected outcomes of the stakeholders' plans defined strategically, to the planned operations pre-tactically, to the actual execution phase, tactically. As future work, the outcome of the downstream layers can be fed back to the previous layers to improve how some decisions are made. This is a typical case of **reinforcement learning, which can be used to make better predictions** and also to optimise the system.

Further scientific questions can be studied by modifying the model. The impact of infrastructure **expansions for airports**, the **way airlines may compete** for new routes, and **new ANSP structures** (including pricing schemes) are but some examples.

More detail on further research and next steps are presented in Section 6.

Acknowledgements

The Vista team thanks all those experts who generously participated anonymously in the various consultation and review exercises, in particular the dedicated workshop in Vienna on 23 October 2017, kindly hosted by Frequentis.

We also thank DLR for kindly providing confidential formulae and valuable advice for the evaluation of emissions impacts, as explained in Section 2.1.1.3.

Thanks are finally extended to the team at SJU for the constructive feedback on various deliverables during the Vista project, in particular on Deliverable 5.1 (the precursor of this final assessment report, to which we occasionally refer regarding final model updates).

1 Introduction

1.1 Objectives of Vista

Vista aims to quantify the current and future (2035, 2050) relationships and trends between a currently non-reconciled set of performance targets in Europe, specifically:

- the trade-off between, and impacts of, ‘regulatory’ and ‘business’ factors;
- the horizontal metric trade-offs *within* any given period;
- the vertical trade-offs *between* periods, particularly as many targets are not currently mapped from year to year, are discontinuous with other targets, or even entirely missing for given periods (such as, vitally, passenger performance targets);
- whether alignment may be expected to improve or deteriorate as we move closer to Flightpath 2050’s timeframe, for example.

Vista focuses on identifying how selected indicators for five stakeholders evolve under different scenarios in the different considered timeframes. The stakeholders considered are: airlines, air navigation service providers, airports, passengers and the environment. Vista will provide insight into the impact of factors on the relationship between the different indicators.

1.2 Overview of Vista model

Figure 1 presents an overview of the Vista model as described in D4.1 (Initial framework definition). Vista models the three temporal phases of ATM (strategic, pre-tactical and tactical) for each scenario investigated, with the objective of generating a representative (busy) day of operations for each given scenario. The different factors define the scenario to be modelled. The strategic layer considers the factors and the economic environment to provide the outcome of strategic decisions made by the stakeholders, the capacities provided and demand, flight schedules and passenger flows for a typical day of operations. These flows, schedules and capacities are transformed into individual flight plans, passenger itineraries and ATFM regulations by the pre-tactical layer. Finally, the tactical layer executes the flights and passenger itineraries at a flight and passenger level, tracking the evolution of delay, passenger connections and the tactical decisions carried out by the actors. Among these layers, only the tactical one is based on a prior model, the Mercury engine. Mercury has been reimplemented and enhanced for Vista, but is based on previous models, used in past projects like POEM and ComplexityCosts. The two other layers are composed of blocks (see Figure 1) which have been written from scratch.

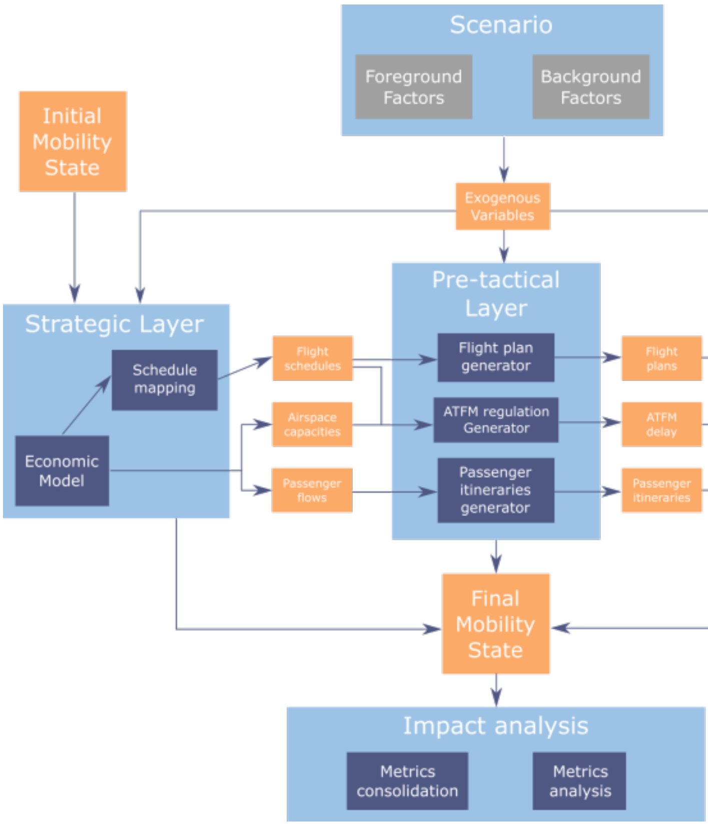

Figure 1. Overview of Vista model

Compared to previous versions of the schema, the feedback loop between the tactical and strategic layer is omitted, without prejudice to the current results. The previous trade-off analysis assessment layer has thus been replaced with a more general 'impact analysis' layer, but has essentially the same function. It represents the set of tools developed by Vista to easily produce trade-off and analysis plots, as presented in Section 5.

One of the capabilities of Vista is to provide metrics for the different stakeholders at a model level, but also at a layer level. **Therefore, it is possible to obtain an indication of how some metrics evolve as the phases move from strategic to tactical.** For example, the ANSPs will estimate their demand strategically and use that information to compute their airspace en-route unit rate. At a pre-tactical level, once the flight plans are generated, the en-route costs planned (to be incurred by airspace users) can be computed, and finally, tactically, once the actual flight plan is selected, the actual en-route airspace charges are computed. Hence, it is possible to assess the evolution of metrics due to the impact of the different layers and the factors.

1.3 Overview of this deliverable

This deliverable presents the final assessment of a set of scenarios. The deliverable presents the:

- details of the modelling for each of the layers (Section 2);
- calibration of the model (Section 3);
- factors modelled and how they combine to create a set of scenarios (Section 4);
- final results obtained with the model, described for the different scenarios and providing insights into the impact of the factors on different metrics (Section 5);
- next steps and future research (Section 6).

1.4 Vista in a few numbers

Vista is an ambitious project, which has compiled high volumes of data to feed a complex multi-layered model. This currently comprises:

- more than 10 distinct sources of data;
- more than 50 GB of data analysed to feed the model;
- a single database with 52 GB of input/output data;
- a GitHub repository with more than 500 commits;
- 93 272 lines of code;
- 67 scenarios simulated;
- around 20 factors simulated.

2 Vista model

2.1 Strategic layer

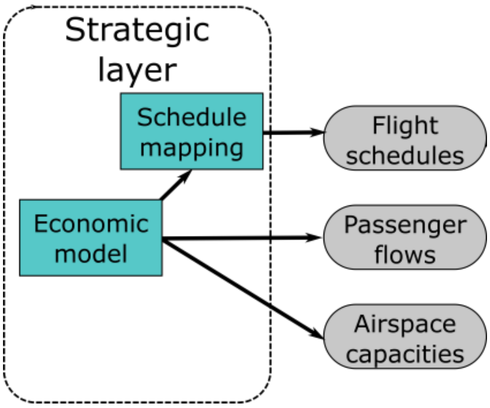

Figure 2. High-level view of the strategic layer

The strategic layer of Vista is designed to produce the flight schedules, the passenger flows and the airspace capacities that are then used by the pre-tactical layer. It is composed of two blocks: the economic model and the schedule mapper.

2.1.1 Economic Model

The economic model is the first block of the strategic layer. It has the ambitious task of creating appropriate levels of supply and demand in Europe based on different scenarios. The model should provide high-level views on the number of flights and passengers, etc., for each origin-destination pair.

With respect to D5.1, the main change in the behaviours of the agents in the strategic layers are the following:

- airports do not change their capacities by their own decisions;
- airports create uncertainty in addition to their average delays;
- ANSPs fix their capacity, taking into account some uncertainty in traffic; they also have a maximum capacity that they can set for a given technological level;

- flights compute a cost expectation based on a Normal distribution of delay with mean and variance fixed by the departure and arrival airports;
- fights include the costs of emissions, which may be assigned to CO_2 and NO_x, and, in future, to H_2O and contrails.

2.1.1.1 Choice of the model

The main requirements of the model are that it should capture:

- the main business drivers in terms of cost and revenues for the airlines in particular;
- complex non-linear economic feedback;
- highly heterogeneous behaviours in terms of agents – airlines, airports, etc.;
- network effects due to alliances, code sharing, connecting passengers, etc.;
- main stakeholder behaviours – airports, airlines, passengers, and ANSPs.

In addition, it should not be computationally over-demanding. For all these reasons, we chose to use a deterministic, turn-based agent-based model, underpinned by a network structure. Agent-based models are particularly well suited to represent different competing agents with heterogeneous behaviours. The agents can have arbitrary complex rules and interact heavily with each other, which allows potential strong feedback between them. We chose a deterministic model to avoid multiple runs as an output, which would require more runs performed by the downstream blocks. Turn-based models are also easier to build and control, even if some scientific issues can arise. For instance, the sequence of 'play', even for the same types of players, sometimes changes the outcome of the game.

2.1.1.2 Structure of the model

The model is based on a common environment, called 'world', in which different agents evolve. The following types of agents are implemented in the model:

- Alliance: one per current real alliance, with some partnerships taken into account too (see Section 3.1.1 for more details). Airlines without any alliance in reality have their own alliances (in which they are alone) in the model, for consistency. These are passive agents, used to build itineraries between airlines of the same alliance. They are not described further here.
- Airline: one per real airline. Mainly used to compute handle the different flights and decide if new OD pairs should be opened.
- Flight: one per OD pair (without connection) per airline. Notionally represents all the flights operated by an airline between the OD pair. Computes the marginal profit of the leg and chooses the supply (number of seats) for the next turn. Also chooses between the possible flight plans based on their costs in fuel and ATC charges. See under 'Flights', below, for how competition is handled.

- ANSPs: one per real ANSP. In the current scenario, set its capacity (number of controllers) based on a target delay. Then sets its unit rate to have zero profit.
- Passenger: one per initial origin and final destination (**with connections**). Notionally represents all the passengers going from 'O' to 'D' by any legitimate itineraries (sequence of airlines and of airports).

Note that there are no hard-coded archetypes of agents within each type. What defines the different behaviours of the agents is their cost structure and their initial conditions (initial network for airlines for instance). As noted above, a (dynamic) network is underlying the structure of the model, where the airports are nodes and flights are edges. On this network, passengers use collection of edges from their initial origin to their final destination.

The supply and demand interacting on this network in an intricate way. On one hand, the supply is leg-based. To compare the supply and demand, the latter is aggregated based on all the passengers going through this leg for this airline, for whatever itinerary. (An "itinerary" is a distinct permutation of legs between a given OD pair, operated by a given alliance.) The supply for the leg for this airline is simply given by the corresponding flight agent. A price variable attached to this leg plays its adjusting role, as we will see in the following. On the other hand, the demand reacts only to price of itineraries. These prices are thus aggregated from each leg in the corresponding itinerary.

2.1.1.3 Agent actions

Each turn, the agents perform a number of actions depending on their nature. In brief:

- based on the scenario, the parameters of the agents are changed;
- ANSPs predict the traffic for this round; they compute the capacity needed to be under a target delay per flight; they set this value for their capacity if it is under their maximum capacity, otherwise they set their capacity to their maximum capacity; they set their unit rate so that their profit is null;
- airports do not change their capacity;
- airlines do not compute whether they should open new routes;
- flights predict the delay at airports, the delay at ANSPs and the price on their leg for this turn; they set their supply based on these predictions and their own cost function;
- passengers take the price of last round for each leg (naive prediction) and sum them to form the price for the itineraries; passengers weight each itinerary based on their utility function (see below) and set their demand level for each of them based on this weighting;
- demand is aggregated for each leg;
- supply and demand are compared for each leg; price evolves on each leg based on the discrepancy;
- all agents record variables value for this round, i.e. prices, delays, etc;

- a new turn is initiated;

In the following, we highlight some details about each agent and their behaviours.

(i) Flights

Flights provide a supply level based on the expected price on their leg. They do this by estimating the price on their leg in the next turn and computing their operational cost. Each flight has a collection of stylised flight plans, a succession of ANSPs with corresponding distances

To compute its operational cost, a flight considers all the possible flight plans between its arrival and destination and weight them to have an average cost. To do this, it collects:

- $\mathrm{p}(\delta \mathrm{t_O})$ and $\mathrm{P_O}$, respectively the probability distribution of delay and the aeronautical charges of the departure airport, from the corresponding agent;
- $p(\delta t_D)$, and P_D, respectively the probability distribution of delay and the aeronautical charges of the arrival airport, from the corresponding agent;
- $\mathrm{p}(\delta \mathrm{t}^{\mathrm{j}}_{\mathrm{ATFM}})$ and $\mathrm{u_j}$, respectively the probability of ATFM delay generated by each ANSP crossed by the flight plan, and its unit rates, from the ANSP agents;
- the fuel price $\mathrm{p_f}$, from the 'world';
- the emission charge per kg of CO_2, p_e, from the world;
- the equivalent CO_2 emissions of NO_x, H_2O, and contrails, e_{NOx}, e_{H2O}, and e_{con}, respectively, for this OD pair;
- the conversion factor between a kilogram of fuel and a kilogram of CO_2 produced, is set to 3.15 in the model (value taken from [1]; IATA uses a very similar value of 3.16).

It then computes the following cost for the flight plan i, with total distance $\mathrm{d_i}$:

$$\mathrm{c}_0^{\mathrm{i}} = E_O[f_d] + E_D[f_d] + \sum_{\mathrm{j} \in \mathrm{ANSPs}} E_j[f_d] + \left(\mathrm{a} + \mathrm{bd_i} + \mathrm{cd_i^2}\right)\mathrm{p_f} + \sum_{\mathrm{j} \in \mathrm{ANSPs}} \left(\mathrm{u_j} \frac{\mathrm{d_i^j}}{100} \sqrt{\frac{\mathrm{MTOW}}{50}} \right) + \mathrm{P_O} + \mathrm{P_D} + (\mathrm{c_{crew}} + \mathrm{c_{maint}})\mathrm{d_i} + 3.15\, p_e\, (1 + e_{NOx} + e_{H2O} + e_{con}) \left(\mathrm{a} + \mathrm{bd_i} + \mathrm{cd_i^2}\right)$$

where f_d is the cost of delay function of the airline, of which the expectation needs to be taken from the delay probability distributions corresponding to the origin airport, to the destination airport, and to each of the ANSPs crossed by the flight plan. In other words:

$$E_O[f_d] = \int f_d(\delta \mathrm{t}) p_O(\delta \mathrm{t}) d\delta \mathrm{t},$$

and so on for the other terms. Expectations such as these need in general to be integrated numerically, which is computationally heavy. However, depending on the shape of f_d and the shape of the probability distribution, some analytical expressions can sometimes be derived. The cost of delay function, taken from [7], is:

$$f_d(\delta t) = \left(d_1 + d_2\sqrt{\text{MTOW}}\right)\delta t + \left(d_3 + d_4\sqrt{\text{MTOW}}\right)\delta t^2 \text{ if } \delta t > 0,$$
$$= 0 \text{ otherwise}$$

This shows two important non-linearities: the quadratic term in the delay and the cut at zero delay. These non-linearities imply that the expectation of the cost of delay should be significantly different from the cost of the expected delay. Indeed, airlines in the project pointed out how important non-linearity of the cost function and unpredictability the operations are for them. In particular, it is clear that an airport creating zero delay on average will still trigger a non-zero cost for the airlines. Similarly, it is clear that high delays are proportionally much more important for the airlines than smaller ones. Using only averages underestimates vastly the cost of delay for airlines.

The fact that the cost of delay function has a simple polynomial form implies that an analytical solution could be found. This would particularly be the case if the distribution of delays were Normal. In the calibration, we show how a Normal distribution fits the data quite well, even though a log-Normal fits slightly better. Since having a closed form for the expectation with the latter is not possible, we decided to use a Normal distribution. If the probability of delay is $\mathrm{p(x) \sim N(\mu, \sigma)}$, then the expectation can be proven to be:

$$E[f_d] = d_l(\mu A + \sigma^2 B) + d_q\left((\mu^2 + \sigma^2)A + \mu\, \sigma^2 B\right),$$

Where:

- $d_l = \left(d_1 + d_2\sqrt{\text{MTOW}}\right)$ is the linear term of the cost of delay;
- $d_q = \left(d_3 + d_4\sqrt{\text{MTOW}}\right)$ is the quadratic term of the cost of delay;
- $A = 1 - \int_{-\infty}^{0} p(x)dx$ is related to the cumulative Normal distribution and thus easily computable;
- $B = p(0)$ is the value of the distribution in 0, also easily computable.

This formula is used for the first two terms of the operational cost equation, i.e. the one corresponding to airports. For the third term, the one corresponding to ANSPs, we used only the cost of the expected delay, i.e. the distribution is a Dirac distribution centred on the mean delay. This is because we do not include uncertainty created by ANSPs (see paragraph (iv) for more details).

The other parameters a, b and c in the cost are the coefficients (specific to the flight) giving the consumption of fuel as a function of the total distance. MTOW is the maximum take-off weight of the flight, $\mathrm{d_i^j}$ is the distance flown in ANSP j airspace, $\mathrm{c_{crew}}$ and $\mathrm{c_{maint}}$ are the cost for crew and maintenance, computed by kilometre flown. Note that the parameters $e_{NOx}, e_{h2O}, e_{con}$ are set to 0 in some scenarios where there are no (current or anticipated) charges for their respective emissions. Where charges are applicable, the model computes these parameters based on a formula kindly provided by DLR, developed from previous research [14]. This formula cannot currently be disclosed for confidentiality reasons. However, the model used to derive the formula is based on detailed propulsion, aircraft design, atmospheric data and so on. The coefficients, dependent on flight length and average OD latitude, are averages. They would be more accurate for the current operational situation than for future ones. Note also that the impact measured is a long-term one, called “ART100”, for “averaged temperature response over 100 years”. This is derived by considering the cumulative radiative forcing effect of the respective chemical component over 100 years. In summary, for a given OD pair, emissions are calculated in terms of equivalent CO_2 emissions, by estimating long-term greenhouse effects with specific coefficients for NO_x, H_2O, and contrails.

Once the cost of each potential flight has been computed, the flight weights the flight plans according to a logit rule:

$$p_i = \frac{e^{-c_o^i/s}}{\sum_j e^{-c_o^j/s}},$$

where s is a smoothing parameter to take into account other possible criteria of choice. This is a free parameter of the model, which drives the 'competition' between flight plans, and thus ANSPs. It is not currently calibrated but will be in the future model, where at least a sensitivity analysis will be performed. These weights can be thought as the probability to choose the given flight plan i, and are in fact exactly used like this in the flight plan generator block. Here in the economic model, they are used to compute the expected revenues for the ANSPs and the expected cost for the flight as a whole:

$$c_o = \sum_i p_i c_o^i$$

Once the operational cost for this turn is known, the flight makes a prediction for price in the next turn with an exponentially weighted memory function, computing the price at round t, p_t, as function of the past N values:

$$p_t = p_{t-1} + \left(\sum_{n=1}^{N} \gamma^i (p_{t-n} - p_{t-n-1})\right) / \left(\sum_{n=1}^{N} \gamma^i\right)$$

The parameter γ tunes the importance of past changes with respect to last one. This is a free parameter of the model. In theory it does not change the final state of equilibrium of the model, only the transient states. Thus, it does not require a full calibration.

This is a fairly standard setup in this kind of model. The flight does exactly the same thing with the delay at departing and arrival airport. Note that this behaviour is required to avoid having unwanted oscillations in the model. Indeed, it is well known that 'naive' variable expectations (i.e. predicting that the price in the next round will be the same than in the last one) leads to unstable systems.

Once expected revenues and operational costs are known, the supply is set by using the following cost function:

$$c(S) \ = \ c_o S + \frac{1}{2} c_c S^2$$

The revenues of the flight being pS, the optimal supply level is easily derived and is given by:

$$S = \frac{p - c_o}{c_c}$$

The supply is in terms of number of seats available and is a continuous variable. The shape of the cost function and its implication are discussed at the end of this section.

Note that no number of flights per se is computed by the flight agent. Instead, all of its internal decisions depend on the number of seats S. It is sometimes required by other parts of the model to compute a number of flights during the day on an OD, notably with respect to the passengers to compute frequencies of itineraries. This is done by considering that the average number of seats per flight is constant on the OD. Hence, this number is computed in the initial state and used throughout the simulations. The subsequent model, the schedule mapper, allows for different aircraft to be used in the schedules.

(ii) Airlines

In the final version of the model, airlines are passive agents that do not formally take any decisions. These are taken at the flight level, using information from the airlines for decision-making.

(iii) Airports

Airports are driven by their cost and delay function. Indeed, following the literature, the airport is assumed to produce a given average delay based on a given traffic delay, with the following linear equation:

$$\overline{\delta t} = \delta t_0 + \frac{T}{C}$$

Where C and δt_0 can be thought of as the capacity and the delay when traffic is very low. Both parameters are directly fitted from data*. The linear relationship between delay and traffic is a simple one which is commonly used in literature and justified with queueing theory. Given a system which can 'process' an entity with a rate μ and under the assumption of a Poisson distribution for arrival entities in the system with a rate λ (M/D/1 system type), the average waiting time (if $\lambda < \mu$) is given by $w = \lambda / \big(2\mu(\mu - \lambda)\big)$. For airports, μ corresponds to their instantaneous capacity (maximum movements per unit of time), and λ represents the traffic per unit of time. Typically, μ is much bigger than λ (over a day for instance), and thus the equation is approximated to $w = \lambda/(2\mu^2)$, where $2\mu^2$ plays the role of effective capacity in our model. Note that the assumption of Poisson-distributed arrivals is a very strong one, but in practice we found that a linear relationship fits the data quite well for most big airports*.

Airports also create uncertainty, i.e. not only a mean delay. This is then used by flights to compute an expected cost of delay, as explained above. We consider the distribution of delays created by the airport as a Normal distribution with mean delay δt and standard deviation σ: the 'uncertainty'. The choice of a Normal distribution is mainly driven by efficiency reasons. Indeed, the expectation of a polynomial function (like the cost of delay function used by the flight) on a Normal distribution has an analytical expression. The remaining equation to add is the one linking the uncertainty to the evolution of traffic. To do this, we used a phenomenological law extracted from data between the mean delay and the standard deviation of the distribution*. This law is a simple linear one:

$$\sigma = \sigma_{oo} + \sigma_s \overline{\delta t}\,.$$

During a simulation step, airports start by aggregating their departure and arrival traffic. Based on that, they compute their delay, which will be then accessed by the airlines to compute costs. Note that, strategically, the delay is assumed to be the same at departure and arrival. The airports also predict traffic in the future and whether or not they should expand their capacity. They do this by comparing their revenues (aeronautical charges plus passenger revenues) to the **operating** cost for their extended capacity, using the cost function:

$$c(C) = c_0 + c_1 C$$

We chose a linear function for this given that there is very little data on the production function of airport. It has already been argued in [3] that a linear law is the best estimate one can make without further data. The parameters in the equation can however be calibrated on real data for some airports. If the profit is high enough, the capacity is extended by a fixed increment after a fixed period of time, both of which are chosen during calibration. Note that, based on the feedback from the workshop Vista organised in October 2017, it has been underlined that most airports cannot extend their capacity, and it has been suggested to have a shortlist of airports which can extend their capacity. We have searched for more information that would help to build such a list. A good candidate was [2], but this document was released too late to be used in Vista, and too much of it is undisclosed with regard to airport capacities. Note that capacities still increase in the model because of external business factors, but not because of internal decisions made by the airports.

Note the aeronautical charges at airports are fixed throughout the simulation, mainly because airports are under a wide range of rules concerning their charges and we lack the data to properly take this into account*. These charges are considered to be the same at departure and arrival.

(iv) ANSPs

ANSPs have been updated based on the feedback from consultation, informal discussions, and important feedback from Belgocontrol. In particular, the ANSPs were previously able to reach their target delay under all conditions. We have now introduced a maximum capacity and changed the way the efficiency is used.

The ANSPs are similar to airports, in the sense that they produce some delay for flights, based on their level of traffic. However, modifying their capacity is much easier for most of them (when they are not saturated already). ANSPs are, however, under current legislation, supposed to keep the average delay per flight under a certain threshold, that in the following we call ‘target delay’, which should be understood as the ‘target maximum delay’. In the simulation, the target delay is a constant and the same for all ANSPs. The delay produced by ANSPs given their capacity and a given level of traffic is given by an equation similar to the airports’, but taking into account the uncertainty in terms of traffic. In a given time window, a given ANSP can only have a maximum of C flights under its control, where C is the capacity, considered fixed throughout the day. Each flight above the capacity will on average be delayed to the next window. Hence, even if in theory the ANSP could choose its capacity with respect to the planned traffic, delays still arise from uncertainty, i.e. unforeseen fluctuations in the number of flights in the area. Belgocontrol has stressed this as being crucial for ANSPs in general.

If p(x) represents the distribution of traffic given an expected traffic T, the average delay created is:

$$\delta = \tau \frac{1}{C} \int_{C}^{+\infty} (x - C) p(x) dx,$$

where τ is the typical magnitude of the delay regulation. Assuming a Normal distribution for the traffic centred in T with a standard deviation σ, the previous expression can be rewritten as:

$$\delta = \tau \frac{1}{C} ((T - C)F(T - C) + \sigma^2 p(T - C)),$$

where p and F are respectively the corresponding centred Normal distribution and cumulative distribution. This expression allows an efficient computation of the mean delay.

This expression reduces naturally in the case where the traffic is completely predictable ($\sigma = 0$), with:

$$\delta = \tau \frac{1}{C} (T - C) H(T - C),$$

where H is the Heaviside step function. In this case, delay is created if and only if the traffic is greater than the capacity, and grows linearly with it, in the same fashion as the equation used for the airports. When σ is increasing, the ANSP has to increase its capacity C in order to keep the level of delay constant. This exactly captures the lost capacity due to unpredictability.

The variance σ is itself a function of the traffic within the ANSPs. Just like with airports, it is expected that higher levels of traffic translate into higher levels of average delay *and* higher levels of uncertainty. In the model, it is computed as the variance of the distribution of delay from the airports. Indeed, we assume that the main factor for the randomness of the traffic levels in the ANSPs' regions is the randomness in the departure delay. Using simple arguments and a Poisson model, it can be shown that:

$$\sigma = \sqrt{T} \left(1 - e^{-\frac{\sigma_\delta}{\tau_0}}\right) \frac{1}{\tau} \left(1 - \frac{1}{\tau}\right),$$

Where:

- T is the average traffic;
- σ_δ is the standard deviation in delay from the airports (which is an average over all the airports);
- τ_0 is the controlling time-window for the ANSP (set to 1 hour);
- τ is the typical delay of a regulated flight (set to 15 minutes).

Note that our definition of the capacity for the ANSPs implicitly assumes that the structure of the traffic inside the airspace does not drastically change between scenarios, i.e. that the complexity is constant.

Given the target delay and having predicted a level of traffic (in the same way that the flights predict price), one needs to invert this equation in order to choose the capacity which is required to be under the delay target. This cannot be done analytically and is performed with a scalar minimisation algorithm using the Brent method. The result represents the ideal capacity that the ANSP should have to keep the delay under the target. This ideal capacity however cannot be always reached, because the ANSP has a maximum capacity. The maximum capacity represents the capacity that is achievable by the ANSP with a constant traffic structure and a given technology. Following discussions between Belgocontrol and UoW, it has been set to 120% of the current capacity for every ANSP. This is obviously a simplification, since in particular different ANSPs are already differently close to saturation. This maximum capacity depends only on the technology, and changes with various factors, in particular the SESAR-related ones.

Once its capacity has been fixed, the ideal capacity or its maximum one, whichever is smaller, the ANSP computes the corresponding cost c and fixes the unit rate u, in order to have zero profit:

$$u = 100\frac{c}{V}$$

Where V is the term of revenue corresponding to $\sum_{\mathrm{i}} d_i\sqrt{MTOW_i/50}$, with the sum computed over all flights crossing the ANSP's region. The parameter V is also predicted by the ANSP independently for this round using the same procedure as for traffic.

Note that like most traffic and capacities in the model, the ANSPs' capacities are given as a number of flights per day.

(v) Passengers

Each agent represents all passengers travelling from an origin and going to a final destination, with all the possible itineraries, including direct routes, connecting flights with the same airlines, and connecting flights with different airlines within the same alliance. The passengers get the price p_k of their itineraries by adding the price of the individual legs (this point is discussed further in the last paragraph of this section). They also compute the frequency f_k of the itinerary by finding the minimum frequency among the legs involved. Based on the different scenarios, they can also have their level of income i_k changed. This impact their utilities, which translates into changing demand. For a given itinerary k, the first factor in the change of demand is given by:

$$\frac{\delta v_k}{v_k} = -\alpha\frac{\delta p_k}{p_k} + \beta\frac{\delta i_k}{i_k} + \zeta\frac{\delta f_k}{f_k}$$

With $\alpha\ is$ the demand elasticity w.r.t price, β is the demand elasticity w.r.t to income, and ζ is the demand elasticity w.r.t to frequency. This corresponds to a 'volume' term v_i, i.e. a term dependent only on the changes of characteristics of the given itinerary. However, passengers are also sensitive to difference of prices between itineraries. To capture this competition, we use the common discrete-choice model based on a linear function:

$$w_k = 1 - \frac{\delta u_k - \sum_i \delta u_i}{s},$$

Where $\delta u_k = -\alpha\delta p_k + \zeta\delta f_k$ is the term of the previous equation under competition pressure (contrary to income). Note that in the literature, the logit function (with exponential) is more common than this one, but exponentials are sometimes problematic computationally, which is why we decided to use a linear function. The final demand for an itinerary is thus $v_i w_i$.

The choice is a simple linear law for the volume term because it is customary in economics to compute elasticities, and thus the calibration was able to benefit from published literature on the subject. Note that if mathematically this can always be justified through a first-order development around a point of equilibrium, in reality, finite changes lead to non-linear behaviours which are not captured by the present model. In particular, the demand is obviously capped by the number of people able to make such trips.

2.1.1.4 Price adjustment mechanisms

Price adjustment mechanisms are independent of any agent formally in the model. Once the supply for a given leg S_i has been computed and the demand D_i for the same leg has been aggregated, the price on leg i undergoes the following change:

$$p_t^i = p_{t+1}^i \left(1 + 2\lambda \frac{D_i - S_i}{D_i + S_i}\right)$$

Where λ is a 'friction' parameter implemented to avoid unstable oscillations in the model. These real, applied prices are then passed back to the agents (pax and flights) for their prediction in the next round.

2.1.1.5 Discussion: price aggregation, network, and supply

Many hypotheses have been made for this model and whilst most of them are discussed above or in the calibration section*, we focus on a couple of central hypotheses in this section.

It is clear that production (in the sense of how many seats will be offered by a given airline on a given leg) is set at an airline level, whereas in the model it is set at the 'flight', i.e. the leg, level. This is done intentionally to avoid a number of issues and keep the model simple enough for our present needs. Indeed, typically given price p, an airline is able to produce a number of goods (or units of service here) S, but with increasing marginal cost:

$$c(S) = S^\alpha$$

In this equation, α is greater than one, which means that each additional produced unit is more expensive that the last one. This type of function is typical in economics and well-established. If now the airline wants to know how many units it should produce given that its revenues will be Sp, one needs to derive the profit with respect to S, which gives *in fine*:

$$S = \left(\frac{p}{\alpha}\right)^{\frac{1}{\alpha - 1}}$$

This is equivalent our airline's supply function presented above, with an added linear term and an exponent $\alpha = 2$. Interestingly, S is an increasing function of p, the price. Given that in general the demand is a decreasing function of the price (it is the case for instance in our model), there is usually a price of equilibrium for which the supply and demand match. Note that in this framework, the market is supposed to be perfect, i.e. the price is not impacted by the supply of a given airline, i.e. the airline has no market power.

We drew on these ideas for our model, with several conceptual changes. We consider indeed than individual legs produce each independently their services (number of seats). We still derive a production function, not using a market price as before, but the price to which they think they can sell (predicted price above). This is some kind of 'micro-market', where each airline is notionally producing a similar, but non-interchangeable, good on each leg. The supply on each leg for different airlines are not aggregated, and thus there is one price per good (per airline, per leg). The airlines are however still in competition through the fact that the passengers can choose between similar different goods (itineraries).

The production function is not at the airline level because this is problematic in non-perfect, concurrent markets. Indeed, given that the airlines optimise their profit overall **for a given price**, their will tend to put all their production in the most profitable leg. The price drops on this leg, which then attracts demand, depleting the other legs which then change their price, and so one. In short, the model is completely unstable in this case. It is important to understand that this is due to the fact that airlines are able to compute changes of demand w.r.t price in the model, i.e. the elasticities. This requires much smarter agents and is thus out of the scope of the present project.

An issue with the current model is that it cannot capture the so-called network effect where airlines are supposedly 'losing money' on one flight in order to feed others profitably. We argue that this vision is deeply flawed, since the profit associated with a flight should be computed by the difference between the case where it is here and the case where it is not. If the airline makes more profit in the case where the flight is operated, then it is arguably the definition of a profitable flight. As a consequence, if the airlines compute the profit for each flight correctly (i.e. taking into account the network effect), in reality they should not have any loss-making flights, by definition.

The canonical way to compute the profit of a single flight is to divide the price of a ticket with a connection, sometimes using weights, for instance based on the distances travelled on each leg. We argue that this is a very crude way of computing the revenues from passengers and does not match any reality. In fact, it is important to understand that the ticket prices cannot be in general disaggregated without involving very arbitrary rules. The only way to compute the revenues or marginal revenues is to take into account the network and its state with and without the flight. This is why the model is based on aggregation of price from legs to itineraries, and not the contrary. The rule of aggregation could arguably be changed. However, it is clear that it should not be far from a simple sum, and in any case the results should not be larger than the sum, because otherwise travellers could just buy their tickets independently to game the system.

* For further information, see Section 3.1 on calibration.

2.1.2 Schedule mapper

The schedule mapper is the second block in the strategic layer. Relatively simple conceptually, it should convert the high-level flows of the economic model into individual schedules, to be used by the flight plan generator. If the role of the schedule mapper is perfectly defined, designing it required a lot of effort due to the high computational complexity of the problem at hand.

Indeed, plan schedules based on expected demand is a highly demanding task, even at a single airline level, who in general has a dedicated tool for that. The complexity of assigning schedules is due to the high number of possibilities and the multiple constraints. These constraints include hard constraints like crew, aircraft, and airport slots and soft constraints like the cost of operating an OD pair.

It is out of scope of Vista to reassign completely from scratch all the schedules created by the economic model. In particular this would imply capturing very complicated slot management behaviours from the airlines (including irrational behaviour like endowment). Moreover, the number flights in play is simply too large to hope for a rather quick stable solution, computationally speaking. As a consequence, Vista simplified the problem by relying heavily on an initial state, set to be 12SEP14 as for the rest of the model.

In short, the schedule mapper compares for each airline the new flow to the old one for each OD. If the new flows are high enough compared to the old ones, each airline tries to add a new aircraft and optimise its route to that most of the new flow is covered. If the new flows are small enough, it removes one of its aircraft to avoid wasted capacity. The airline takes crew and airport slots into account only indirectly through the possible patterns (route) available to the new aircraft and the corresponding likely turnaround times.

More specifically, the mapper goes through the following steps:

- get data on airports, historical schedules, pattern data (see next subsection), and the strategic flows;
- compute average travelling times between every OD pair, even those which are not present in the historical schedules;
- compute the likely departure times;
- get the decision tree for the turnaround times;
- for each airline:
    - trim its network by removing aircraft which in excess;
    - grow the network by adding aircraft to meet demand;
- compute the new schedules and add them to the database.

In the following subsection we explain in more detail the different concepts behind this procedure.

With respect to the earlier model version (presented in D5.1), the schedule mapper includes the following changes:

- a simplification has been made concerning the patterns; each airline now uses only back and forth for their aircraft, dismissing all the triangular, quadrangular, etc. patterns; the code is much faster with this new setup;

- some further heavy machine learning techniques have been used to choose turnaround times; a decision tree is now used by airlines to do this, based on a limited number of predictors;
- new aircraft are now chosen based and capacity range.

2.1.2.1 The concepts of taxon and pattern

The schedule mapper relies heavily the successive steps that an aircraft follows during a given period of time and how the aircraft of an airline together deliver the desired capacity.

We use the name 'pattern' as the labelled succession of airports that a given aircraft makes during a given period of time, for instance 'LHR, CDG, LHR, ORY'. Patterns can be of two types: either a 'loop', i.e. the first and the last airport are the same, or an open pattern. Note that a pattern does not include times of departure, travelling times, etc.

Going one step further, we define a 'taxon' as an unlabelled succession of airports that a given aircraft makes during a given period of time. For instance, the patterns 'LHR, CDG, LHR, ORY' and 'FRA, MAD, FRA, FCO' both share the same taxon, which can be coded for example as '1, 2, 1, 3'. Hence, from a given taxon, many patterns can be derived, from which even more final schedules can be generated (adding departure times).

This allows us to produce a hierarchical algorithm, in which the airline first needs to find an adequate taxon, among a relatively small number of them.

The possible taxons and patterns are extracted from data, as described in Section 3.1.1.3. This section in particular shows how a very small number of taxons is enough to capture most of the airlines' network.

Among these taxons, some can be said to be 'elementary', in the sense that one can generate all the possible taxons by combining them. Driven by the analysis in Section 3.1.1.3, Figure 3 shows the three elementary taxons considered in the analysis. They represent respectively an aircraft going back and forth between two airports, an aircraft going successively through three airports, and an aircraft going successively through 4 distinct airports. All of them are looped, i.e. they cannot generate open patterns.

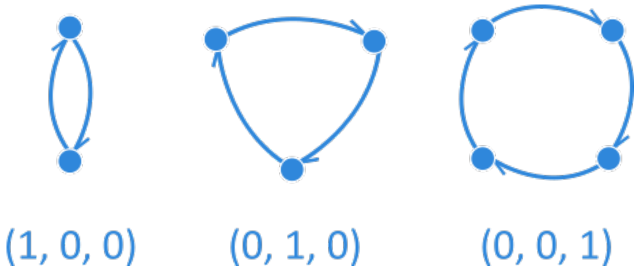

Figure 3. Elementary taxons

Each circle is a distinct airport. Each arrow represents an aircraft going from one airport to another.

These elementary elements can generate a lot more taxons by combining them in different manners. Strikingly, only a few of them are significantly present in the data. Figure 4 shows most of them (for a

72 hour time window), and the corresponding notation. For instance, the taxon on the left is a repetition on the elementary taxon (1, 0, 0) -> (2, 0, 0). Note that we assumed that the same airport is always the start of the elementary taxon (i.e. we excluded taxon such as '1, 2, 3, 4, 3, 2, 1'). The available taxons represent a common pool from which the airlines can draw.

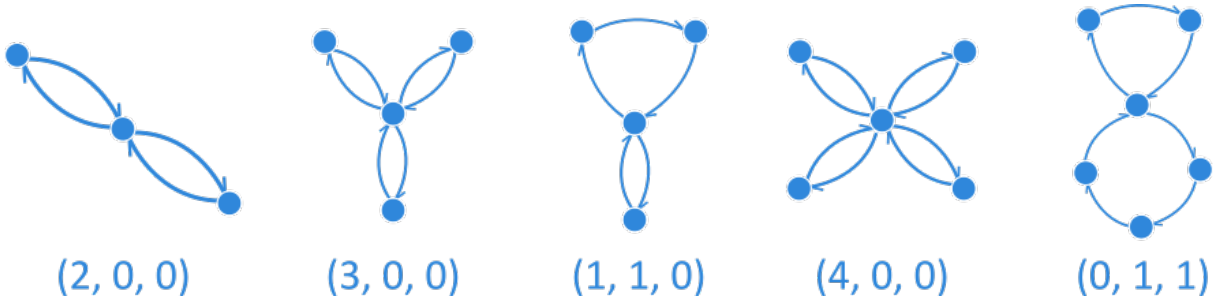

Figure 4. Composite taxons

The airports are not necessarily unique in this representation.
For instance, the taxon on the left can represent two back and forth trajectories between the same two airports.

Since the vast majority of taxons used by airlines are composed of elementary 'back and forth' ones (see the calibration section), we used only these types in the final model (in contrast to previous versions).

2.1.2.2 Airline cost function

Each airline tests whether or not they should remove current patterns or add some. In order to take a decision, the airline uses a 'master' cost function, used for every pattern. It is built by computing the vector $\mathrm{DD} = \mathrm{D} - \mathrm{S}$, representing the difference between demand and supply on each leg, and the corresponding vector of price P. We then restrict these two vectors to the elements where $\mathrm{DD_i} > 0$, i.e. where there is more demand than the supply, and note them DD^* and P^*. The cost is then:

$$\mathrm{c} = \mathrm{c_c N_p} + \mathrm{DD}^*.\mathrm{P}^*$$

The first term is a cost of capital, where $\mathrm{c_c}$ is the cost for creating a pattern, loosely corresponding to the cost of an aircraft (buying or renting). $\mathrm{N_p}$ is the number of patterns. The second term corresponds to a shortfall: it is the money the airline would gain by increasing the supply to meet all the demand.

2.1.2.3 Removing under-used aircraft

Removing aircraft from the network allows us to better prepare the next stage of adding aircraft to maximise aircraft utilisation. To do this, the airline accesses the current patterns and tries to remove them to see if the situation improves, based on the above cost function. We used for this a simulated annealing technique, with Boolean variables corresponding to each existing pattern. In short, the algorithm tries to switch randomly each variable on and off and computes the difference of cost between the two situations. It goes mostly towards minimum cost, with some probability to increase the cost, based on a parameter called 'temperature'. The temperature is slowly decreased to find the absolute minimum in cost. Once this procedure is complete, the corresponding optimal solution is passed to the next step, where new patterns are added.

2.1.2.4 Taxon and pattern optimisation

The procedure takes advantage of the fact that aircraft need to ultimately come back to their origin airport after one (or more) flights. We only consider patterns composed of 'back and forth' elements.

Using this fact, the airline first considers which would be the best OD to provide a flight for. This is done by finding the indices i and j for which DD_{ij} is maximum, i.e. the OD where there is a maximum discrepancy between demand and supply. Considering only this edge, the airline then chooses the aircraft which would be best suited for the connection, using its own historical aircraft plus two generic aircraft (B747 and A320). Indeed, the airline checks whether an aircraft has a range compatible with the OD pair, and takes a generic one, otherwise. The departure time from the origin is then chosen from historical ones.

Once the first edge has been chosen, the same aircraft operates the return flight. To find the departure time from the destination (which is the new origin), it picks a turnaround time using a decision tree, based on the following variables:

- the previous time of arrival;
- whether or not the airport is in the ECAC space;
- which type of airline is operating the flight;
- the number of passengers on the flight;
- the distance flown during the previous flight.

Using the tree, the airline chooses randomly one of the possible options for the turnaround times, which fixes the next departure time. More details on the tree can be found in the calibration Section 3.1.2.

After using the same procedure to find the next departure time, the procedure then finds the next best edge starting from the same airport, and so on. This greedy algorithm stops when a flight is finally outside the time window or when the estimated gains are negative, using the master cost function. This pattern is saved and another one is initiated by finding the next best edge (not necessarily from the same airport).

2.2 Pre-tactical layer

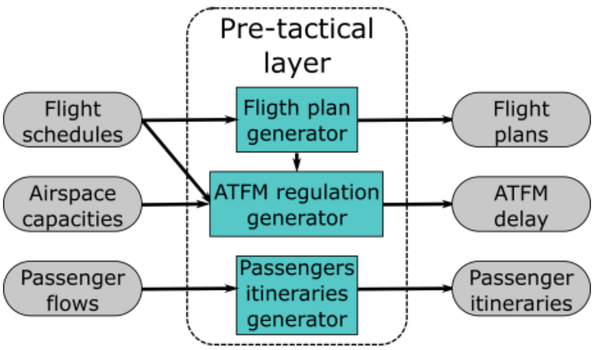

Figure 5. High level pre-tactical layer blocks

As shown in Figure 5, the objective of the pre-tactical layer is to create the necessary level of detail to execute tactically the flights and passengers' itineraries. The outcome of this layer are individual flight plans, ATFM regulations and probabilities of delay being assigned due to a regulation and individual passengers itineraries.

The layer is divided into three blocks:

- flight plan generator;
- ATFM regulation generator;
- itineraries generator.

2.2.1 Flight plan generator

The flight plan generator transforms origin-destination schedules into defined possible flight plans. For each flight plan an estimated operating cost is computed including fuel, CRCO and delay/buffer costs. These operating costs are taken into consideration when prioritising the different flight plans for each schedule.

This block is based on historical possible trajectories and on aircraft performances (based on observed trajectories and on aircraft performances models (BADA)).

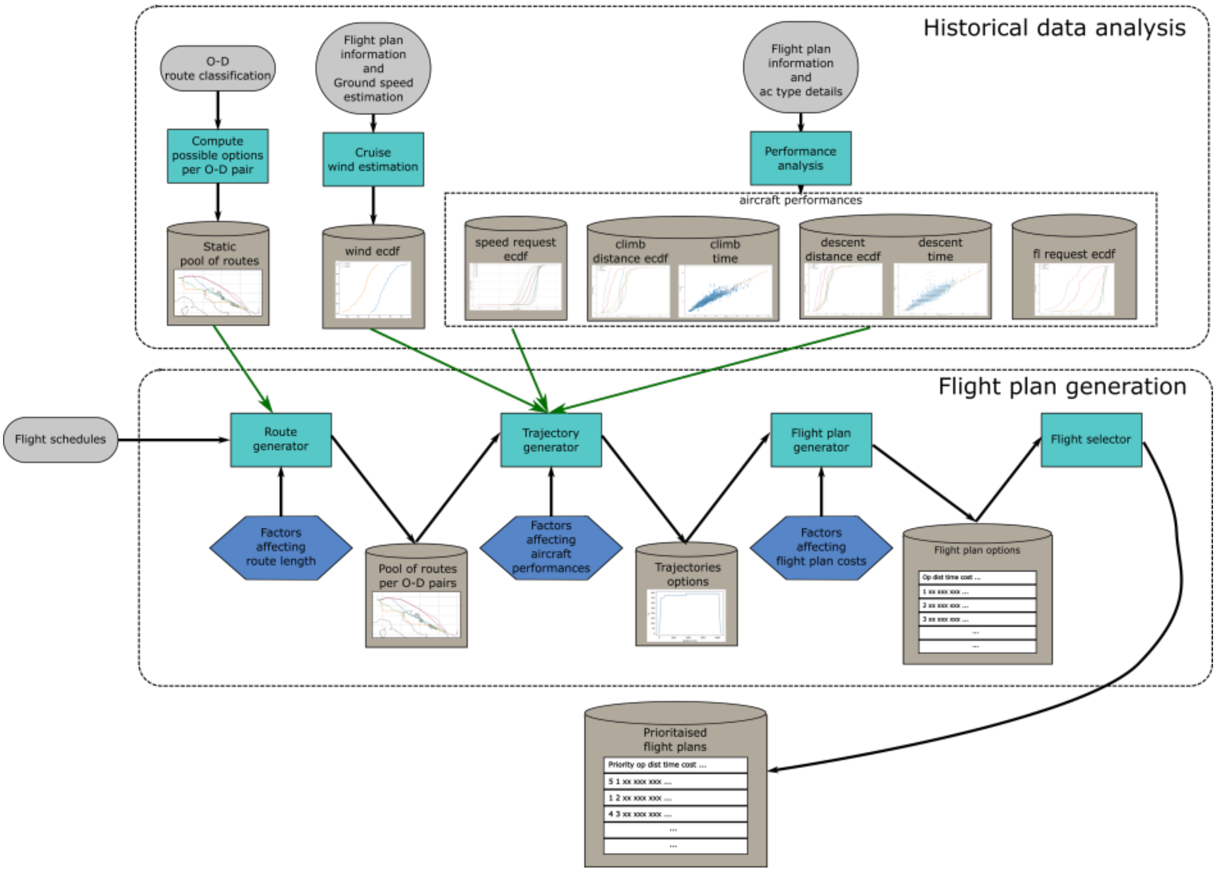

Figure 6. High level pre-tactical layer blocks

The flight plan generator is divided in different sub-blocks, these blocks are supported by historical data analysis on flight trajectories and aircraft performances (see Figure 6). These historical and performance data analyses are described in more detail in Section 3.2.1.

2.2.1.1 Route generator

For each scheduled origin-destination pair a set of possible routes are computed. These possible routes are based on the historical possible routes flown between these origins and destinations (see Section 3.2.1.2 for details on how these options are computed from historical data). If the origin-destination pair does not exist in the historical dataset a set of new routes are estimated as follows.

All the routes between the NAS[1] of the airports of origin and destination are considered. The routes selected by the generator are those which minimise the distance between the airport of origin and destination and the origin and destination of the routes. This mean selecting the routes i which minimises the following equation:

$$min_i(GDC(O, O_i) + GDC(D, D_i))$$

O and D are the airports of origin and destination and O_i and D_i the airports of route i. In many cases one of the airports (either the origin or the destination) will have a route that links it to an airport in the NAS of the other airport. Finally, for each NAS where the airport in the routes selected are not the airport of origin and destination, these are linked to the entry/exit points of the route in the NAS considering the great circle distance. This distance is increased by an inefficiency factor computed as a function of the difference between the entry/exit point in the NAS and the location of the airports.

[1] Note that if there are no flights between a given origin and destination, all flights between their NAS of origin and destination are considered.

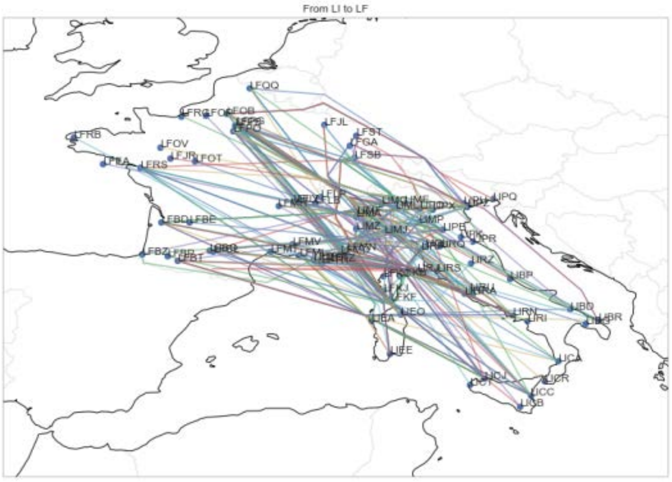

a) All possible routes between LI and LF

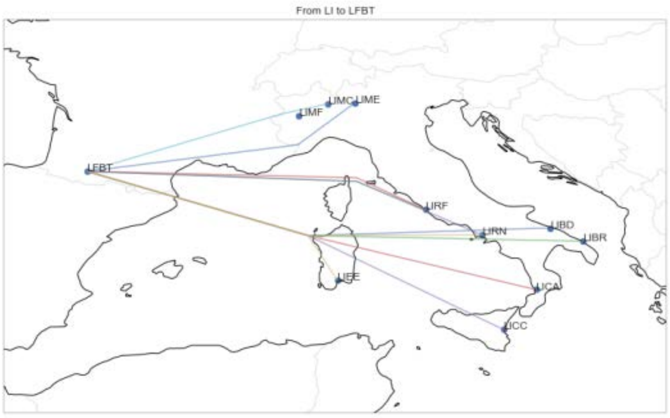

b) Routes between LIMF and LF

c) Routes between LI and LFBT

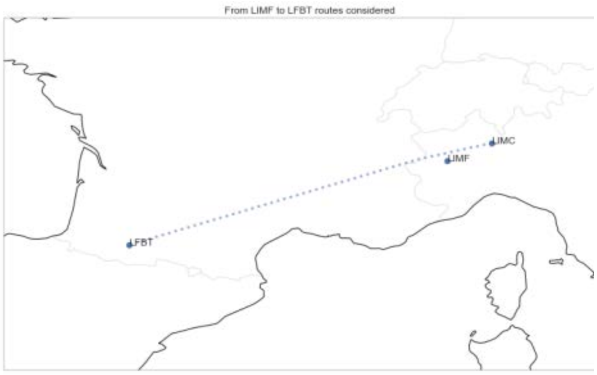

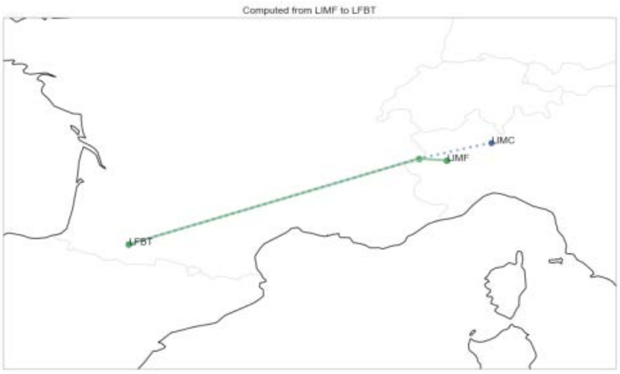

d) Routes that minimise distance to LIMF and LFBT (LIMC – LFBT)

e) Route generated LIMG – LFBT

Figure 7. Example of route generation LIMF – LFBT

Figure 7 shows an example of this route generation between LIMF and LFBT. In this case the route that minimises the distance between its origin and destination and LIMF and LFBT is an existing route between LIMC and LFBT. The route is modified as shown in Figure 7 e) by linking LIMF with the exit point in the Italian airspace of the LIMF to LFBT route.

If there are no flights in the pool of historical flights between the origin and destination NAS of the origin and destination, then the possibility of routing via a third airport is analysed. Figure 8 shows an example for flights between MUHA and LSZH.

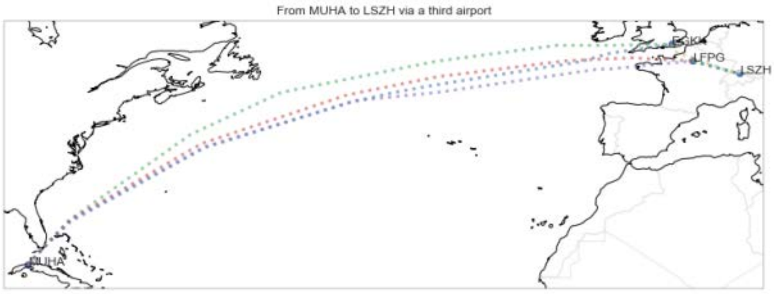

a) Options selected for MUHA – LSZH via EGKK and LFPG

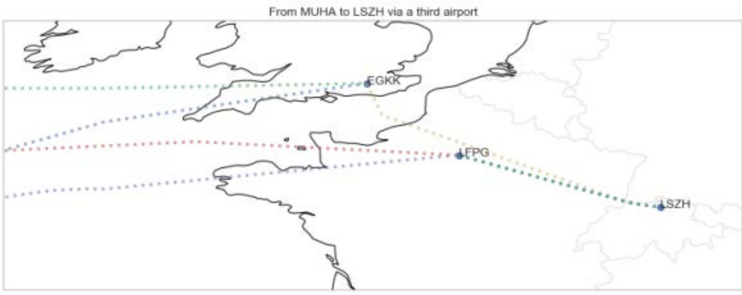

b) Detail of MUHA – EGKK – LSZH and MUHA – LFPG – LSZH

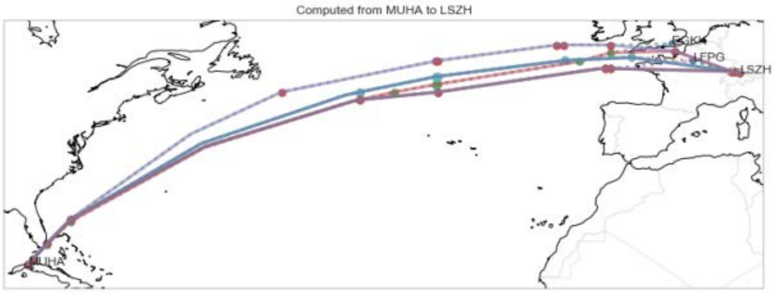

c) Routes generated MUHA – LSZH

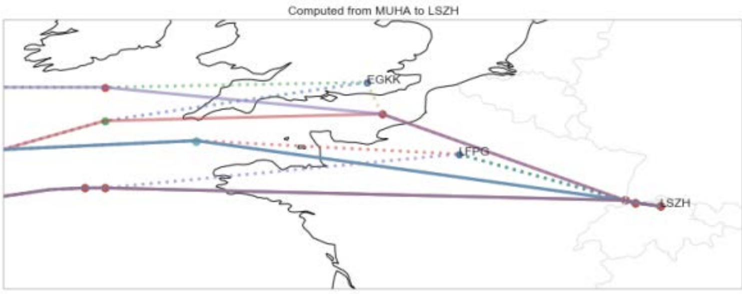

d) Detail of routes generated MUHA – LSZH

Figure 8. Example of route generation MUHA – LSZH

In this case, the shortest route between MUHA and LSZH via a third airport is via LFPG. All the intermediate airports which have a route linking MUHA with LSZH which is shortest that the longest route MUHA-LFPG-LSZH are selected as potential intermediate airports to compute the new routes. In this example, as shown in Figure 8 a) and Figure 8 b), the airports selected are LFPG and EGKK. This removes options such as MUHA-EDDF-LSZH or MUHA-UUEE-LSZH which are possible but unrealistically long. Finally, in the NAS of the intermediate airport, the entry point of the MUHA – intermediate airport is linked to the exit point of the intermediate airport – LSZH route, as presented in Figure 8 d). The distance within this NAS is increased from the great circle distance considering an inefficiency factor computed based on the extra distance between the entry point and the intermediary airport and between the intermediary airport and the exit point with respect to their GCD as shown in the following equation:

$$\textit{inefficiency factor} \quad = \frac{GDC(Entry\ point\ NAS, airport) + GDC(airport, Exit\ point\ NAS)}{Distance\ (Entry\ point\ NAS, airport) + Distance\ (airport, Exit\ point\ NAS)}$$

By generating these routes, the route generator will ensure that for each origin and destination in the schedules there is at least one possible route in the pool of routes.

The characteristics of the routes within the pool of historical and estimated possible routes will be modified based on the factors. In particular, the length of each segment within each NAS is reduced based on the background factors of SESAR improvements considering the variation due to the enhancement of the Environment (i.e., the reduction of the flight distance with respect to the GCD as defined by the SESAR targets).

2.2.1.2 Trajectory generator

The trajectory generator creates, for each schedule, a set of possible trajectories (routes and flight plan characteristics (e.g., climb, cruise and descent distances and times, estimated fuel usage, average cruise wind). The outcome of this block are flight plans: for each schedule a flight trajectory is estimated per possible route between the schedule origin and destination pair.

A trajectory optimiser could have been used to generate the different flight plans. However, with that approach idealised optimal trajectories would have been obtained. Instead, in Vista, the trajectory generator relies on historical data on flight plans and aircraft performances. See Section 3.2.1.2 for more details on the inputs required for this block.

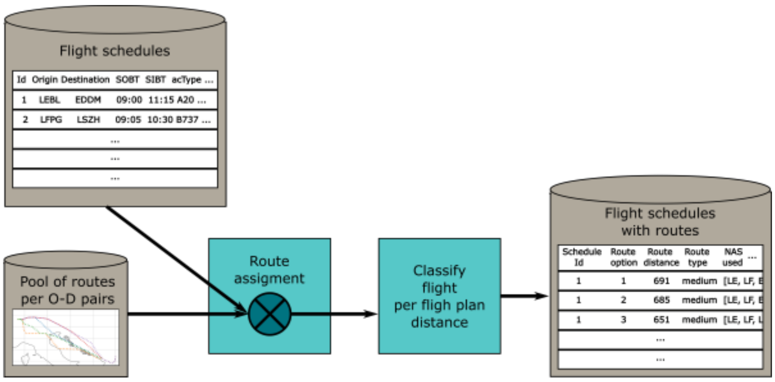

Figure 9. From flight schedules to routes per flight schedule

The first step to generate the trajectories is to join the schedules with the pool of possible routes between airports. As shown in Figure 9, for each schedule all the possible routes linking its origin and destination are selected. These routes are classified based on their length into categorical labels (short-, medium-, long-, very-long-haul flights). This categorisation is used when generating the trajectory characteristics.

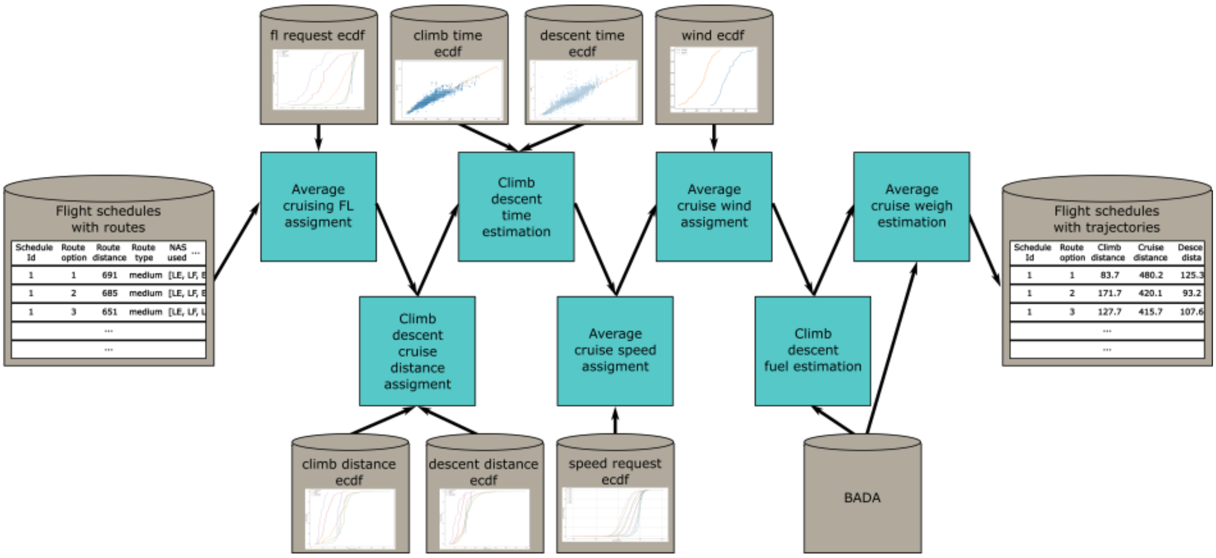

Figure 10. From routes to trajectories

Figure 10 presents the steps required to generate the trajectory for each of the possible routes. First an average cruising FL is assigned based on distributions of flight level requested per aircraft type and flight plan length. A climb and descent distance is estimated considering the same parameters: aircraft type and flight plan distance, considering this and knowing the total flight plan distance, the cruise distance can be computed.

Next, using climb and descent distance a climbing and descending time is estimated. This is done by using linear fittings of climb and descent times as a function of their distance and adding randomness considering the variance of the actual times with respect to the fitting.

An average cruise speed request is estimated considering the distribution of speed requested as a function of cruising flight level. This speed request is done ensuring that the value is within the aircraft limitations.

For each origin and destination NAS an average cruise wind distribution has been estimated (see Section 3.2.1.2 for more details). Using these distributions, the average cruise wind is estimated for each trajectory.

Finally, using BADA performances, the required climb and descent fuel are computed. Considering these values along with the cruise distance, flight level, cruising speed, wind and an estimation of the landing weight (currently landing with 80% of the maximum payload) the fuel and the average cruise weight are estimated. Considering the conversion between fuel burn and CO_2, the amount of carbon emissions are also estimated per trajectory.

In this manner a full flight plan trajectory is generated per route. As there are elements in the process that are stochastic, this block can be executed n times if required to obtain different possible trajectories per schedule. However, we have performed one execution per strategic input as the variability is not significant between executions and all the trajectories generated are realistic.

2.2.1.3 Flight plan generator

The flight plan generator block computes for each trajectory their estimated direct operating cost. For each route, the CRCO charges are computed and a cost of fuel is used to translate the expected fuel consumption into a total expected cost.

EUROCONTROL's central route charges office (CRCO) charges airspace users for the air traffic services used by airspace users on behalf of EUROCONTROL's member states. The charges are computed according to the following formula:

$$CRCO_{cost} = \sum_{n} u_n \times \left(\frac{d}{100}\right) \times \sqrt{\frac{MTOW}{50}}$$

With n being each of the different charging regions, u_n the unit rage in cents of euro and d the great circle distance between the entry and the exit point of the charging region (if the aircraft departs or arrives to the region 20 km are subtracted).

In Vista, these charges have been modelled for the 39 charging regions managed by the CRCO plus the airspaces of Egypt, Belarus, Morocco, Uzbekistan and Ukraine. Other surrounding countries which might follow different charging schemes are also modelled:

- Algeria, which uses the same formula to compute the charges but has different unit rates depending if the traffic is overflying or travelling to-from their country.
- Iceland, which for traffic to-from Iceland follows the same formula, but for traffic crossing the Atlantic applies a fixed rate of GBP 41.14 plus a factor that depends on the great circle distance but not on the aircraft MTOW.
- Russia, which uses the great circle distance but different unit rates depending on the aircraft MTOW defined by categories.
- Tunisia, which have a rate as a function of the aircraft MTOW.

The communication charges of Shanwick airspace are also considered in Vista.

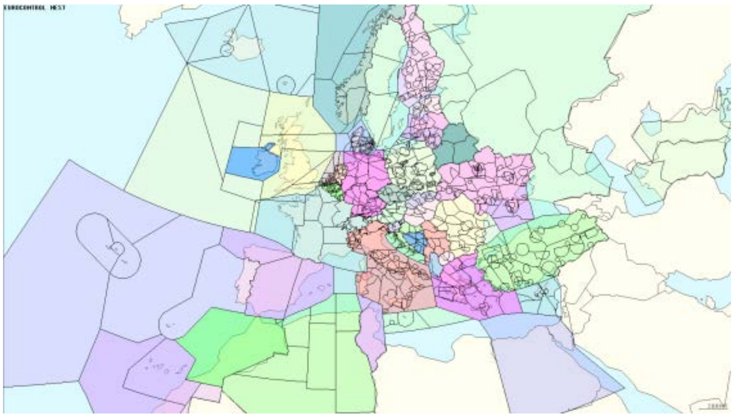

Figure 11. Regions for which the air navigation charges are computed

By including all these regions, the air service charges of the European airspace are covered including surrounding countries, as shown in Figure 11 (note that there are no routes overflying Libyan airspace). This allows us to compare the cost of different routes even when they use adjacent airspaces to the core European ones (for example flying over Algeria to cover a route between the Canary Islands and Italy or using the Shanwick airspace in flights between the UK and the Canary Islands).

The cost of the unit rate used for the core ANSPs considered in Vista are either from historical values when the current scenarios are modelled or the outcome of the economic model of the strategic layer. The value of fuel and carbon emissions also depend on the scenario considered.

2.2.1.4 Flight plan selector

The trajectory generator produces as many trajectories as possible routes between the origin and destination exist. However, the tactical layer requires a flight plan per flight. A flight plan per schedule is also required by the ATFM regulation generator to estimate the demand on each ANSP and, from this, the probability of having an ATFM regulation.

It is well known that airspace users might change their flight plan prior departure as a function of the tactical situation. If a flight has a high ATFM delay assigned a different, longer route that avoids the regulation might be submitted instead; or if a flight is about to propagate delay, a more direct route might be selected. These tactical changes to the flight plan cannot be decided at the pre-tactical layer. For this reason, the Flight Plan Selector will prioritise all the different flight plan options of the different flights based on their operating cost.

Each possible flight plan has a different cost and utility due to differing amounts of fuel, airspace and carbon emissions cost, other parameters could be considered. As shown in Figure 12, for a given flight plan there are buffers at arrival and, knowing the minimum turnaround time and the subsequent SOBT of the aircraft frame, a turnaround buffer time can also be estimated. This leads to the fact that different flight plans will have different probability of being delayed on their own (i.e., having an arrival time after their SIBT) and of propagating delay to subsequent flights (i.e., having an arrival time which uses all the arrival and the turnaround buffer). These considerations could be used when prioritising the different flight plan options.

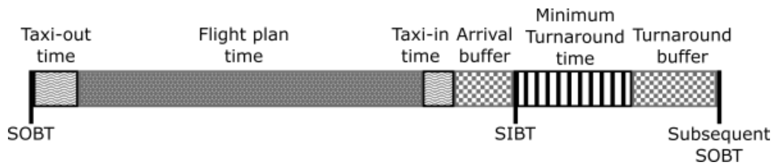

Figure 12. Flight times and buffers

A logit rule, which considers the cost of fuel, the cost of airspace charges and the cost of carbon emissions is used to estimate the probability of selecting a specific flight plan, in a similar manner as explained in Section 2.1.1.3(i) for the strategic layer.

2.2.2 ATFM regulation generator

The ATFM regulation generator estimates the probability of being affected by ATFM regulations and the intensity (amount of delay) due to those regulations. As presented in Figure 5, the input of the ATFM regulation generator is the capacity of the ANSPs (for the scenario to be processed and for the baseline 2014 scenario) and the traffic (demand of the scenario to be processed and of the baseline 2014 case). The ATFM regulation generator thus requires the outcome of the flight plan generator (and flight plan selector) and the 2014 demand and capacity to be able to compute the variation of demand and capacity with respect to the 2014 case, which is used as reference.

ATFM regulations are divided between regulations due to capacity issues (i.e., regulations flagged as “C” as their reason to be implemented by the FMP) and all the other regulations. It is assumed that regulations due to capacity are affected by the demand in the ANSP while the regulations that are not due to capacity remain homogeneous across the ECAC region. This assumption allows us to modify the probability of having ATFM delay due to the demand expected at different ANSPs and their expected capacity, while maintaining delay, which is not directly linked with a capacity/demand imbalances (e.g., regulations due to weather).

Note that from historical data (analysis of AIRAC 1313-1413), as shown in Table 1, regulations due to capacity represent 43.3% of all the regulations issued, followed by weather as the main causes. Weather, by its nature, might occur at wider location across the ECAC and might not be related directly to demand.

Table 1. ATFM regulations by reason excluding industrial actions during AIRAC 1313-1413

ATFM reason	Number of regulations	
Capacity	243 740	(43.3%)
Weather	145 172	(25.8%)
Aerodrome capacity	52 405	(9.3%)
ATC staffing	42 283	(7.5%)
Other	79 676	(14.1%)

With this approach, there is a given probability of having delay due to ATFM regulation which is not dependent on the demand and is the same across all the flights. A given experimental cumulative probability distribution function based on the analysis of a year of historical regulations (AIRACs 1313-1413) provides the amount of delay assigned to the flight.

If a flight is not affected by a regulation set for non-capacity demand imbalance issues, then there is a probability of incurring a delay due to an ATFM regulation due to capacity, which depends on the ANSP airspace that the flight crosses. Each ANSP has a given probability of assigning delay, which depends on the time the flight is crossing the ANSP's airspace and on the expected demand and capacity. The analysis of the regulations issued by the different ANSPs allows us to determine that the probability of having a regulation during the day is spread homogeneously. For this reason, we assume a homogenous probability of having delay due to ATFM regulations due to capacity for each ANSP from a given starting time to a given end time of the day. These times have been computed to encompass the times when the demand is high and vary per ANSP as shown in Section 3.2.2.2.

For each ANSP, the entry counts in a 30-minute window between the starting time and the end time when regulations due to capacity might be issued, and are computed as shown in Figure 13. The addition of all the number of flights in each of the 30-minute window period for each ANSP gives us the demand for an ANSP that is generating regulations due to capacity.

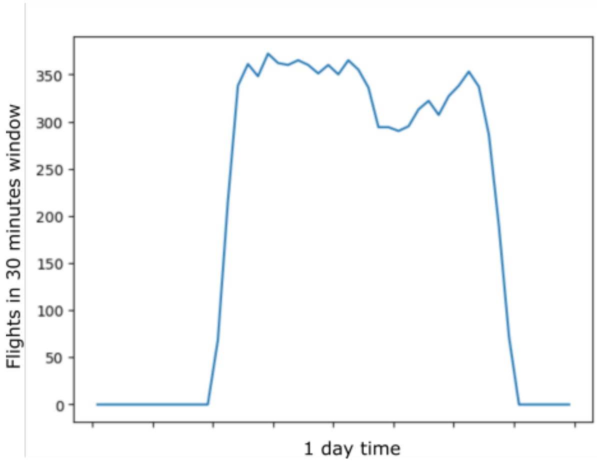

Figure 13. Demand in LE ANSP for 2014 scenario

For the baseline 2014 scenario the probabilities of having delay due to capacity are based on historical analyses of flights affected by regulations in that ANSP, as described in Section 3. The amount of traffic that is estimated in the ANSP for the current scenarios is the original demand (D_0) of each ANSP. If the scenario is not the baseline 2014, then the probabilities of having a delayed flight are adjusted following the equation below:

$$P = P_0\left(\frac{(D_1 - D_0)}{D_0} - \frac{(C_1 - C_0)}{C_0}\right) + P_0$$

D_1 being the demand for the scenario being analysed, D_0 the demand at the 2014 baseline scenario, C_1 the capacity of the ANSP for the scenario being analysed, C_0 the capacity of the ANSP at the 2014 baseline scenario, and P_0 the probability of having delay by crossing the ANSP's airspace due to capacity issues in the 2014 baseline scenario. The capacities of the ANSP (C_0 and C_1) are an outcome of the economic model, and the demand is estimated for the scenario as shown in Figure **13**. As mentioned, there are two experimental cumulative probability distribution functions for the ATFM delay, one for the delay not due to capacity and one for the delay due to capacity.

2.2.3 Passenger itineraries generator

One of the outputs of the strategic layer is the passenger flows. For example, in the 2014 baseline scenario a flow is 878 passengers that would like to go with Lufthansa from EGLL to EDDF. The strategic layer also reports the schedules that the airlines are providing with their aircraft type assigned for each flight leg and the number of seats available on the aircraft.

The objective of the pre-tactical layer is to transform the flows of passengers into individual itineraries, i.e., assign the passengers to specific flights. This is done in a three stage process: computing the possible options available for the passengers in each flow, then optimising the assignment of passengers among their options and finally creating additional passengers' itineraries to ensure that the load factors of the aircraft are realistic and indicating which passengers are 'premium' and which are 'standard' (see Section 2.3.2).

2.2.3.1 Options for passenger flows

If a passenger is making one or several connections, the minimum connecting time required to change between flights is considered when making the computation of possible options. For example, in the 2014 baseline scenario there is a flow with five passengers which goes from Barcelona (LEBL) to Tromsø (ENTC) via Frankfurt (EDDF) and Oslo (ENGM) (LEBL –EDDF – ENGM – ENTC). The following table presents all the possible 18 options for that flow. As observed, there are three departing flights from LEBL which allows the passengers to make all the connections. However, at EDDF there are two possible flights one leaving at 10:20 and one leaving at 14:00. With the 10:20 flight there are five options to keep the route from ENGM to ENTC; and with the 14:00 flight only two of those options are possible as the arrival time plus the minimum connecting time is larger than the departing time of the connecting flights. Each option will have a different amount of waiting/connecting time at the airports, even if in some cases the total waiting time might be the same, since the final flight is the same (i.e. we are sometimes only shifting the waiting from one airport to another). As presented, if departing later from LEBL, fewer options are available.

Table 2. Possible flight options for flow LEBL – EDDF – ENGM – ENTC (all times in UTC)

Option		SOBT1	SIBT1		SOBT2	SIBT2		SOBT3	SIBT3	Waiting time at EDDF	Waiting time at ENGM	Total waiting time
1	F1	04:55	07:05	F4	10:20	12:20	F6	14:20	16:10	3:15	2:00	5:15
2	F1	04:55	07:05	F4	10:20	12:20	F7	15:35	17:25	3:15	3:15	6:30
3	F1	04:55	07:05	F4	10:20	12:20	F8	16:35	18:25	3:15	4:15	7:30
4	F1	04:55	07:05	F4	10:20	12:20	F9	17:55	19:45	3:15	5:35	8:50
5	F1	04:55	07:05	F4	10:20	12:20	F10	20:05	21:55	3:15	7:45	11:00
6	F1	04:55	07:05	F5	14:00	16:00	F9	17:55	19:45	6:55	1:55	8:50
7	F1	04:55	07:05	F5	14:00	16:00	F10	20:05	21:55	6:55	4:05	11:00
8	F2	06:05	08:15	F4	10:20	12:20	F6	14:20	16:10	2:05	2:00	4:05
9	F2	06:05	08:15	F4	10:20	12:20	F7	15:35	17:25	2:05	3:15	5:20
10	F2	06:05	08:15	F4	10:20	12:20	F8	16:35	18:25	2:05	4:15	6:20
11	F2	06:05	08:15	F4	10:20	12:20	F9	17:55	19:45	2:05	5:35	7:40
12	F2	06:05	08:15	F4	10:20	12:20	F10	20:05	21:55	2:05	7:45	9:50
13	F2	06:05	08:15	F5	14:00	16:00	F9	17:55	19:45	5:45	1:55	7:40
14	F2	06:05	08:15	F5	14:00	16:00	F10	20:05	21:55	5:45	4:05	9:50
15	F3	8:30	10:40	F5	14:00	16:00	F9	17:55	19:45	3:20	1:55	5:15
16	F3	8:30	10:40	F5	14:00	16:00	F10	20:05	21:55	3:20	4:05	7:25
17	F4	10:45	12:55	F5	14:00	16:00	F9	17:55	19:45	1:05	1:55	3:00
18	F4	10:45	12:55	F5	14:00	16:00	F10	20:05	21:55	1:05	4:05	5:10

2.2.3.2 Passenger assignment optimisation

Now that all the options are computed we need to assign passengers to these options, deciding how many of the five passengers are assigned to each of the different options. Note that this is a complex problem as we need to ensure that the flights' capacities are respected (i.e., do not exceed the maximum capacity of the aircraft) and that there are interactions between flows. For example, if in this case three passengers are assigned to option 18, then there will be three more passengers in F4, F5 and F10, so one assignment is affecting three aircraft, and there might be other flows that also could use those flights.

To solve this, an integer programming problem has been formulated as shown below:

With:

$$F_{fo} = 1 \; if \; flight \; f \; is \; used \; in \; option \; o$$
$$I_{io} = 1 \; if \; itinerary \; i \; uses \; option \; o$$
$$I_o = number \; of \; passengers \; assigned \; to \; option \; o$$
$$C_f = capacity \; flight \; f$$
$$V_i = volume \; of \; passengers \; in \; flow \; i$$

The objective is:

$$max\left(\sum\nolimits_{o} I_o\right)$$

Subject to:

$$\sum\nolimits_{o,f} I_o \times F_{fo} \leq C_f$$
$$\sum\nolimits_{o,i} I_o \times I_{io} \leq V_i$$

The objective is to maximise the number of passengers that are assigned to an option while maintaining the number of passengers on the flights lower than the capacity of the flight and not assigning more passengers from an itinerary flow than the volume of that given flow.

For each aircraft, its capacity is defined as a target load factor computed as a triangular distribution between 0.75, 0.95 and 1 of their maximum number of seats. This ensures that flights are not systematically full up to their maximum capacity obtaining load factors that are realistic.

If the optimisation is run as defined above has the risk that the optimiser might prioritise single leg itineraries, as they are easier to accommodate and to maximise the number of passengers allocated. However, we want to preserve as much as possible the two- and three-leg itineraries, as those are more relevant in terms of delay propagation through the network and single leg itineraries can be added afterwards as 'fillers' if needed (see below).

Table 3. Percentage assignment different optimisation options scenario Baseline Low 2035

	All together	Only leg 2 itineraries	Only leg 3 itineraries	All together constrained leg 2	All together constrained leg 3	All together constrained leg 2 and leg 3
Passengers assigned	96%	9%	0%	93.7%	95%	95%
Itineraries assigned	88.1%	67%	5.3%	93.7%	88.7%	96.7%
Itineraries with only 1 leg	100%	-	-	99.5%	100%	99.4%
Itineraries with 1 and 2 legs	84.7%	95.9%	-	95.9%	84.8%	95.8%
Itineraries with 1, 2 and 3 legs	77.9%	-	99.5%	37.4%	99.5%	99.8%
Average load factor	84%	14%	0%	83%	84%	83.8%

Table 3 presents the results of assigning the passenger with this optimisation but subject to different constraints. Note that if all the itineraries are optimised at the same time, 96% of the volume of passengers is assigned. 100% could not be obtained due to the constraints on aircraft capacity with respect to those estimated by the strategic layer of the model. Those passengers are from only 88.1% of all the possible itineraries, i.e., 11.9% of the itineraries are not assigned to any option as their flights are already full. All itineraries with one leg are assigned, but only 84.7% of itineraries with two legs and 77.9% of itineraries with three legs are able to be allocated. The average load factor obtained on the aircraft is 84%. These results are good in terms of passengers assigned and load factors, however, they are poor in terms of itineraries allocated, especially for connecting passengers. The table also shows the results if only two-leg or three-leg itineraries are allocated. In these cases, the number of passengers assigned are small (9%, and less than 1%) but this is due to the fact that there are fewer passengers making connections than single-leg passengers. However, we can see that the number of itineraries that can be allocated increases significantly for these connecting passengers, reaching 95.9% of two-leg itineraries being used in the optimisation process and 99.5% of the three-leg itineraries, from 77.9%, having at least one of its passengers assigned when only three-leg itineraries are considered.

Next to be tested was the optimisation of the assignment of the two-leg itineraries, then adding this assignment as a constraint prior to optimising all the itineraries; and likewise for the three-leg itineraries. The results obtained are better in both cases. However, optimising two-leg itineraries and setting them as a constraint still produces a low number of three-leg itineraries to be assigned (only 37.4%) and optimising three-leg itineraries and setting them as constraint reduces the number of itineraries assigned to only 88.7%.

Finally, the option selected is to optimise at the same time the assignment of itineraries with two legs and three legs and to set the outcome of this optimisation as a constraint for the following optimisation of all the itineraries. With this, the penalisation on the number of itineraries optimised is small and 95% of the passengers are assigned. The approach taken in Vista is thus a two-stage optimisation process to obtain the assignment of itineraries from the strategic flows: first an assignment of only connecting passengers (with two or three legs) and then using the outcome of this optimisation as constraints for the optimisation of all the flows.

2.2.3.3 Passenger types and 'fillers'

After the optimisation, some flights might have a load factor which is unrealistically low. This could be due to potential misalignment between passenger flows and schedules and due to constraints on flight capacities. For this reason, a new target load factor occupancy is set on the flights as the maximum between the current load factor and a triangular distribution between 0.35, 0.8 and 0.8. This might mean that more itineraries are needed. These will be generated as 'fillers' of single-leg passengers on those flights. The amount of passengers added as fillers is defined by the load factor of the aircraft, ensuring realistic load factors for the flights. This means that there might be an expected discrepancy with respect to the number of passengers planned at the strategic level, even if this is small as the capacity provided strategically is made considering the expected demand. (The number of filler passengers ranges between 2.5% and 7.1% of the total number of passengers on each scenario, being on average 5.7% of the passengers in a given scenario).

These filler passengers do not have a fare associated to them, therefore a linear fitting is used to estimate the cost of a given fare based on the fare of passengers doing single-leg itineraries and the flight distance. Some randomness is added to this value by computing it as a Normal distribution centred in the linear fitting. Figure 14 shows an example for the baseline low 2035 scenario of this distribution of fares as a function of flight distance.

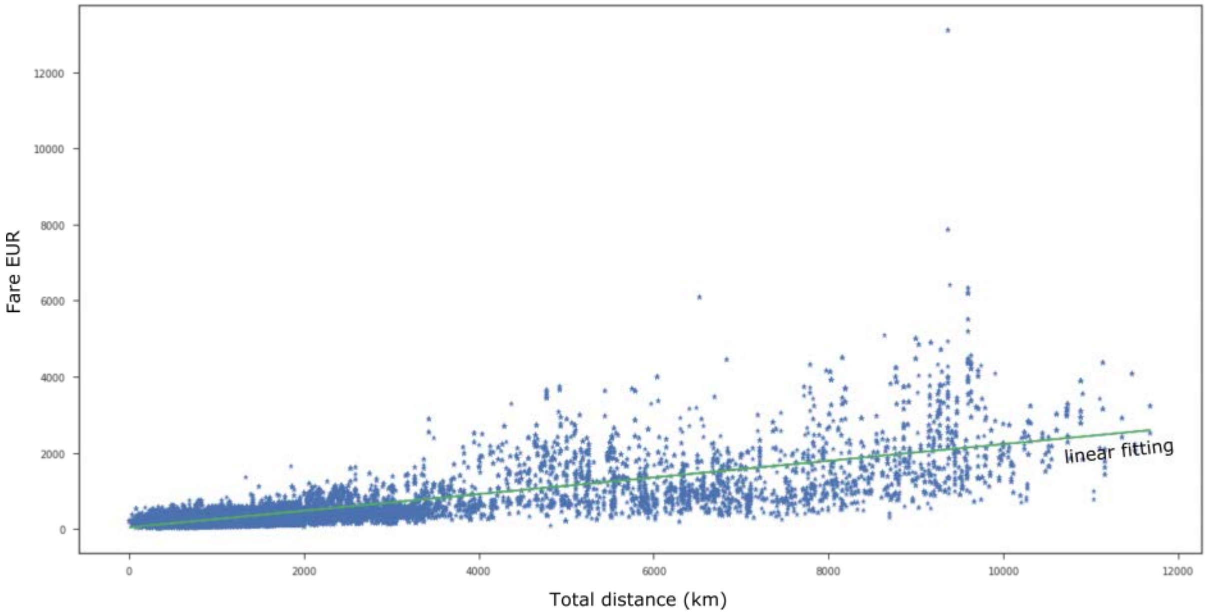

Figure 14. Fare as a function of single-leg itinerary distance for scenario baseline low 2035

Some of the passengers should be premium passengers. A linear fitting of the fare and the total itinerary distance flown (adding the distances of all the legs) has been computed. This fitting gives us a reference of the average price paid and using this we can compute the extra price paid per passenger as the difference between the fitting and their actual value.

All passengers with an extra payment greater than -0.25 [(2)] are considered as *potentially* 'premium' passengers (meaning that in the Mercury model they have reaccommodation priority during disruption; see Section 2.3.2). This is done so as not exclude passengers that have paid less than the average as potentially premium passengers (see Figure 15 for an example of this). The number of passengers that are premium per itinerary is computed as a probability which depends on the extra price.

[2] A dimensionless value: extra relative cost = (actual fare – fitted fare) / fitted fare.

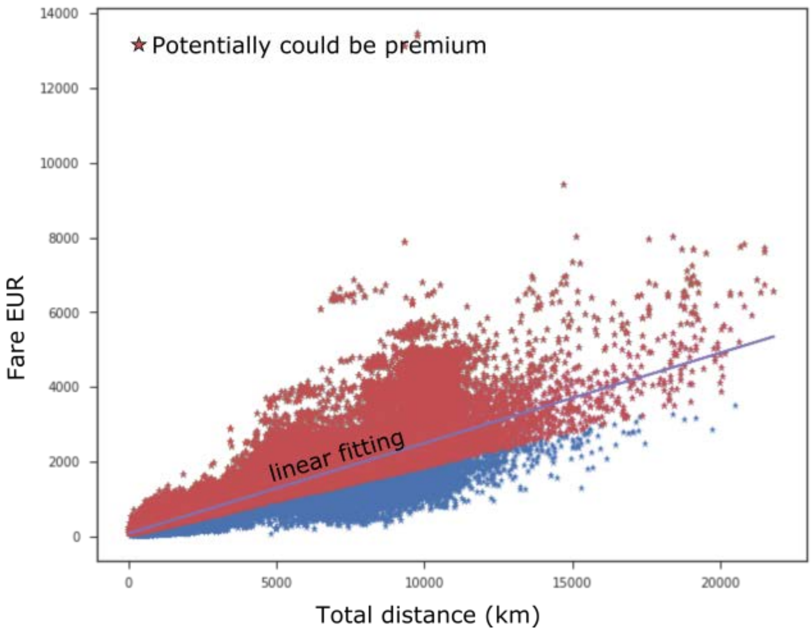

Figure 15. Fare as a function of total itinerary distance and potential premium passengers

Fare as a function of total itinerary distance and potential premium passengers for scenario baseline low 2035.

2.3 Tactical layer

The tactical layer analyses the delay propagation between flights and the adaptability of the system during disruption (cancellations, background and foreground delay) and with limited resources (e.g. airports and en-route capacity). The analysis is performed both per flight and per passenger to allow us to measure the trade-offs between passenger and flight performance.

This layer is integrated in the Vista platform making use of the data and metrics generated in previous layers to simulate the selected scenarios and their underlying factors. However, when provided with the necessary inputs, the tactical layer is independent from the other Vista layers. It is thus possible to replace other layers or to use historical data from other sources.

One of the characteristics of the tactical model proposed for Vista is its door-to-door approach. This tactical layer reproduces the whole passenger journey including the transport choices of the passenger to reach the airport from their origin to their destination. This model was validated during the last ten years across several research projects and it is an evolution of the Mercury simulator developed for Vista. Conversely, the improvements made during the project will be integrated into Mercury.

There are various feedback loops considered in this layer. One example is the arrival delay of flights, which is updated several times during the simulation. Airline operators are affected by such delay because a connecting flight is notified in order to adapt its behaviour to the current status of the system.

Several costs are considered during the simulation, such as fuel costs, crew costs, ANSP costs and passenger delay costs. Those costs have an impact on the behaviour of the system such as when selecting the flight plan. This model also updates the cost metrics considering the last state of the system (i.e. when the simulation ends). Environmental emission costs are also considered, in particular the CO_2 emission costs, as the NO_X emission costs are already modelled in the strategic layer.

The tactical layer simulation is a sequence of two processes, the gate-to-gate simulation and the door-to-door simulation, that are executed for each flight. It requires some information from previous layers, such as the flight schedules, flight plan options, airspace data and passenger itineraries, but also some other platform inputs, such as the aircraft profiles and categorisations. It generates several performance metrics in three database tables in order to be further processed later for creating meaningful visualisations.

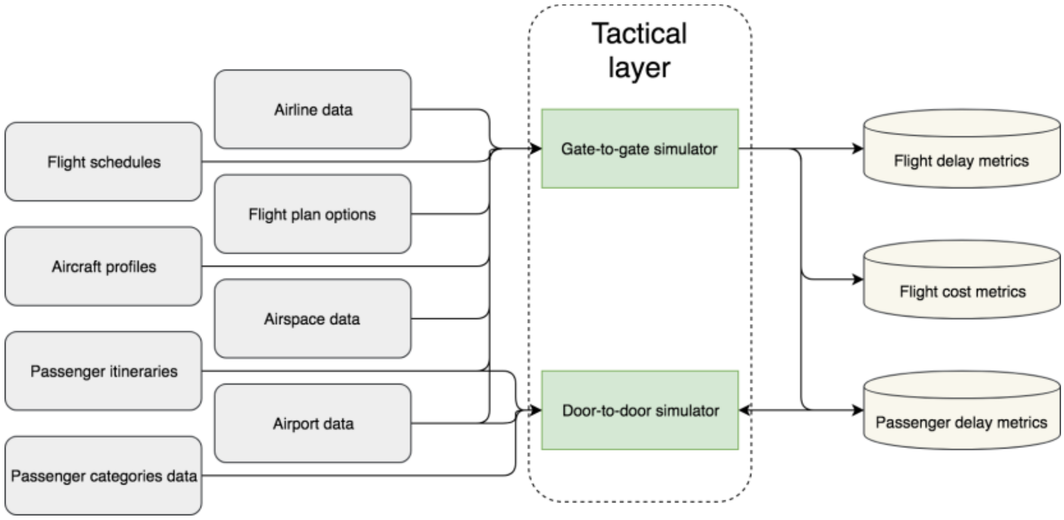

Figure 16. High level tactical layer blocks

A detailed description of each block follows.

2.3.1 Gate-to-gate simulation

The gate-to-gate simulator models each flight and passenger itinerary capturing flight and passenger metrics (e.g. delay, costs, number of missed connections) while modelling pre-tactical and tactical disruptions in the system (i.e. delay due to ATFM regulations and due to other reasons, such as tactical capacity limitations regarding runway occupancy, and uncertainty in the different flight phases).

Instead of simulating the whole continuous domain between the first flight that departs and the last flight that arrives, we have discretised continuous intervals in several discrete events. The time between these events is not fixed: the simulator only evaluates those events that are important for coordination, such as those related to system capacity (runway occupancy, TMA control, en-route control). Once a flight event is processed, it adds the next event to be simulated in the events queue. The event scheduler is responsible for maintaining the event list accessible during the whole simulation and ensuring that all events are processed in order (synchronised). This simulator considers four processes:

- **flight plan submission**: contains all the initial tasks required to select the flight plan to submit as the final flight plan;
- **previous aircraft ready**: contains all the tasks done before updating the flight plan in order to propagate any reactionary delay;
- **push-back-ready**: contains all the tasks done from the push-back to approaching the arrival TMA;
- **arrival**: contains all the tasks done from the arrival TMA to in-block.

Each flight is associated with a unique aircraft, but the same aircraft could be associated with several flights in a sequence. Therefore, four different situations may arise, the flight:

- does not have any previous flight or next flight;
- has a next flight only;
- has both a previous flight and a next flight;
- has a previous flight only.

All these situations are covered in the event chain diagram below. It also identifies all the processes that can be simulated as well as the time events that trigger its execution.

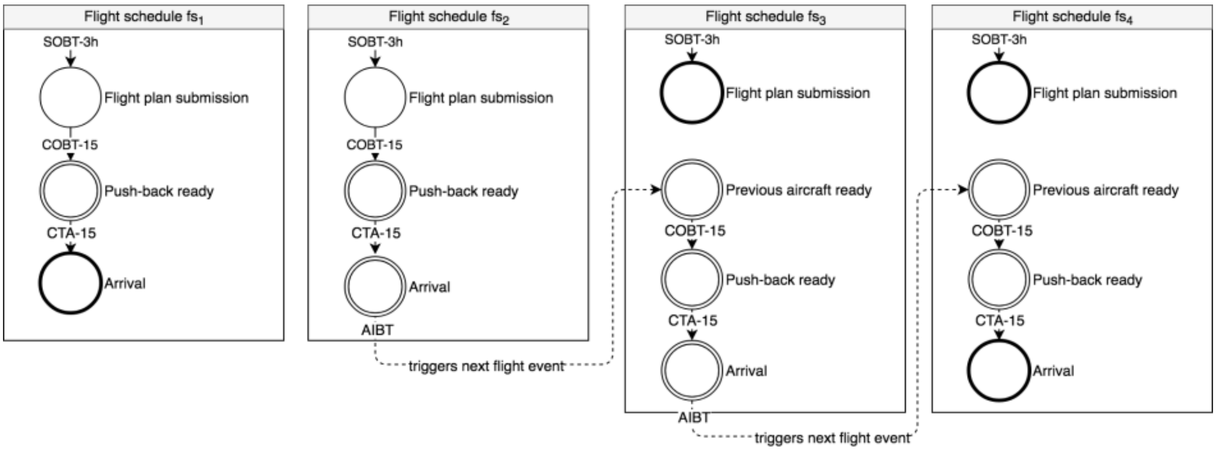

Figure 17. Event chain diagram of processes involved in the gate-to-gate simulation

This simulator has several improvements over the previous version (of D5.1), such as the:

- ATFM model is more realistic and is calibrated with data generated in previous layers;
- en-route model is more realistic as it includes the effect of wind;
- cost model is included.

2.3.1.1 Flight plan submission

Each flight schedule can be either cancelled or confirmed 3 hours before the scheduled time of departure. The confirmation should inform the flight plan that has to be executed later.

The selection of cancelled flights is chosen randomly from probability distributions previously calibrated. If a flight is cancelled:

- affected departures and arrivals boards are updated with the status of the flight; each airport keeps these boards updated to find alternative itineraries when reaccommodating passengers;
- affected passenger itineraries are identified and ‘requested’ a free reaccommodation (as described in the push-back ready section, below);
- the next flight is notified not to wait for this flight (note that the model assumes the next flight with the same aircraft will be offered and no reactionary delay is expected for this cancelled flight).

If a flight is not cancelled: both primary and reactionary delays should be calculated in order to estimate the departure time. This model defines primary delay as a combination of primary ATFM delay and primary non-ATFM delay. The first is calculated after the selection of the flight plan, as an ATFM regulation depends on the airspaces the flight plan crosses.

This module sets the reactionary delay as the difference between the scheduled time of departure and the time the aircraft is actually ready at the gate, based on the previous arrival time and the minimum turnaround time. The turnaround time model is calibrated per airport archetype, aircraft type and airline type. The primary non-ATFM delay is then set taking a sample of delay from a probability distribution.

Next, the flight plan is selected from different options provided by the pre-tactical layer, using reactionary delay and primary non-ATFM delay as a preliminary estimation of departure delay. The model assumes that all the delay is propagated in this phase, so the flight option with the lowest cost will be selected.

Later, the process sets whether the selected flight plan is affected by an ATFM regulation or not. This is done by submitting the flight plan with the estimated departure time to the network manager module described in Section 2.3.1.5. The network manager will provide the delay assigned to the flight plan, if any.

Finally, the departure time is updated taking the new delay estimations into account. Using the updated departure time, it checks all connecting passengers whether they will miss their connection(s). Depending on how many passengers miss their connection(s), the amount of delay they have and if they have an international destination, the departure time could be updated to wait for them, according to the scenario.

The process ends by submitting the final flight plan. It includes an estimation of taxi times collected from the taxi time dataset.

2.3.1.2 Previous aircraft ready

Only flight schedules that have a previous flight with the same aircraft are affected by this process. It is performed as soon as the previous flight aircraft arrives at the airport gate.

This process updates the estimated departure time considering the actual reactionary delay and checking whether the remaining delay has also changed, in this case the ATFM delay, the turnaround time and the non-ATFM delay. The ATFM delay could increase if the flight is already regulated and the new estimation of departure time differs more than 15 minutes. In this case, the ATFM slot will be missed and a new slot request is sent to the network manager, as explained in Section 2.3.1.5.

2.3.1.3 Push-back ready

This process starts 15 minutes before the latest estimation of the departure time for each flight. The preceding process could be either the flight plan submission or the previous aircraft ready process, depending if the flight has a previous flight with the same aircraft or not.

First, this process calculates the taxi-out time by taking a sample from the taxi time model calibrated for each airport. The taxi-out time and the departure time defines the minimum time a flight could be managed before its departure. The departure manager could manage to start the departure process at the minimum time or later, depending on the congestion at the airport. It finally sets the actual take-off time and the departure time is shifted accordingly.

The actual departure time and scheduled departure time is compared and the delay generated is used to calculate the flight departure delay costs, both for maintenance and fuel-related costs. Once the departure time is set, passengers that are not already at the gate are reaccommodated.

Passengers that miss a connection between flights are reaccommodated onto another flight. Firstly, for all passengers for a flight, we calculate the estimated time when those passengers should be at the gate (considering the minimum connection time and the previous leg arrival time) and filter those passengers that arrive out of time at the gate.

The simulator then gets the list of passengers of a flight and removes those passengers from the list in order to (later) try to add more passengers to the flight that might be waiting at the airport, if any.

Next, the flight looks up in the arrivals and departures board of the current airport and the destination airport and gets those flights that connect between that OD pair. This model not only considers direct flights but also indirect ones (up to one stopover).

Finally, the simulator sorts those new itineraries per arrival time and checks if one of those itineraries is suitable for the reaccommodation policy.

If there is no itinerary suitable for free reaccommodation, the passenger could either return to their origin or wait at the airport. If they choose to return, the same procedure described above is done but changing the destination to the origin.

Once all passenger reaccommodations are processed, the simulator offers the empty seats (up to a maximum load) to the passengers that are waiting at the airport.

Finally, the process updates all en-route times considering the effect of wind, the departure delay and the variability of climb time and cruise distance. The updated en-route times are used also for calculating the actual fuel and maintenance costs en-route.

2.3.1.4 Arrival

This process starts 15 minutes before the latest estimation of landing time for each flight. Firstly, the holding time is calculated depending on the current capacity of the airport and the traffic level. The airport capacity model is calibrated using the input from previous layers (in this case, the strategic layer). A positive holding time will update the estimation of the landing time. The actual arrival time at gate depends on the actual taxi-in time, calculated by taking a sample of the taxi time model, previously calibrated for each airport. The actual and scheduled arrival time at gate are compared and the delay generated is used to calculate the flight arrival delay costs.

2.3.1.5 Network manager

As described in Section 2.2.2, the ATFM delay is estimated from two different sources of delay: regulations due to capacity demand imbalances and regulations not due to capacity demand imbalances. The pre-tactical layer has estimated the probability of being regulated due to non-capacity regulations and the probability that each ANSP could give delay to the flight due to capacity-related regulations. Finally, the pre-tactical layer has also estimated the cumulative distributive function of the delay assigned if the flight is affected by a regulation (one distribution for regulations due to capacity and one distribution for regulations due for other reasons).

A module network manager has been implemented to replicate the behaviour of the central entity, which assigns delay to flights if their flight plan is affected by a regulation.

The airline provides to the network manager the flight route with the entry points of the different ANSPs that will be crossed and its estimated climb, cruise and descent distances and times. With this information, the network manager estimates the entry time of the flight for the different ANSPs. If the flight departs from one of the airports that is within an ANSP that form part of the ANSPs that can assign delay due to ATFM (i.e., ECAC, EUROCONTROL and Adjacent States with regulated traffic incoming to the NM area as defined in [18]) then it can potentially have delay assigned.

The network manager will first randomly check if the flight gets regulated by a regulation that is not due to capacity issues. If this is the case, then it will draw the amount of delay using the appropriate cumulative distribution function.

If the flight is not delayed due to non-capacity related regulations, then it could still be regulated due to capacity. In this case, it will select the ANSPs that are expected to be crossed by the flight within the start and end time that each country can regulate the traffic. Each of the potential ANSPs that can regulate the flight is drawn randomly so they are explored in a random manner. The random probability of the first ANSP is used to see if that ANSP is regulating the flight, if this is not the case, the probability of the remaining ANSPs is adjusted, taking into account the ANSP in question, and a new ANSP is randomly selected. This process is done until all the ANSPs are explored, if none of them assigns delay, or until one of them regulates the flight. If the flight is regulated the amount of delay is drawn from an experimental cumulative distribution function.

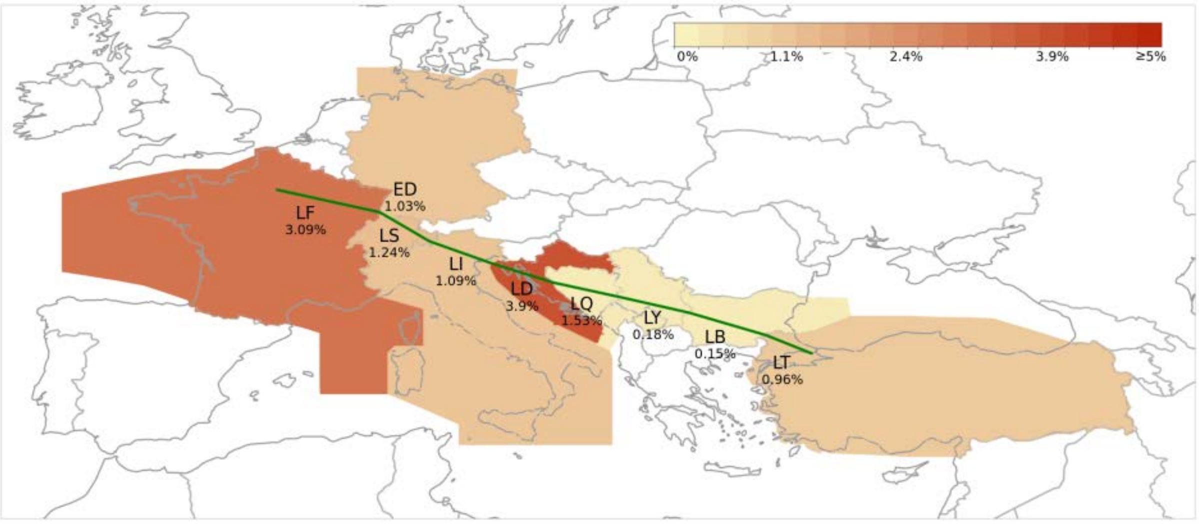

Figure 18. Example of flight from LFPG to LTBA

Airspaces show probability of assigning delay due to capacity related ATFM regulations.

Figure 21 presents an example of a flight crossing several ANSPs from LFPG to LTBA. Each ANSP has a probability of regulating the flight due to capacity issues (e.g., LF 3.09%, ED 1.03%, LS 1.24%). If all the countries can assign delay to the flight, as it is expected to enter them when their potential regulations are active, the network manager would select an ANSP at random (e.g., LI) and test if the flight is regulated (here, there is a 1.09% probability). If the flight is not regulated then a second ANSP will be randomly selected (e.g., LQ) and its probability adjusted considering the fact that the previous ANSP did not regulate the flight. These adjustments are done to ensure that the total amount of traffic regulated per ANSP remains at the probability of each of them due to the conditional probabilities.

2.3.2 Door-to-door simulation

The door-to-door simulator transforms gate-to-gate passenger itineraries into door-to-door itineraries and calculates several passenger metrics for those itineraries.

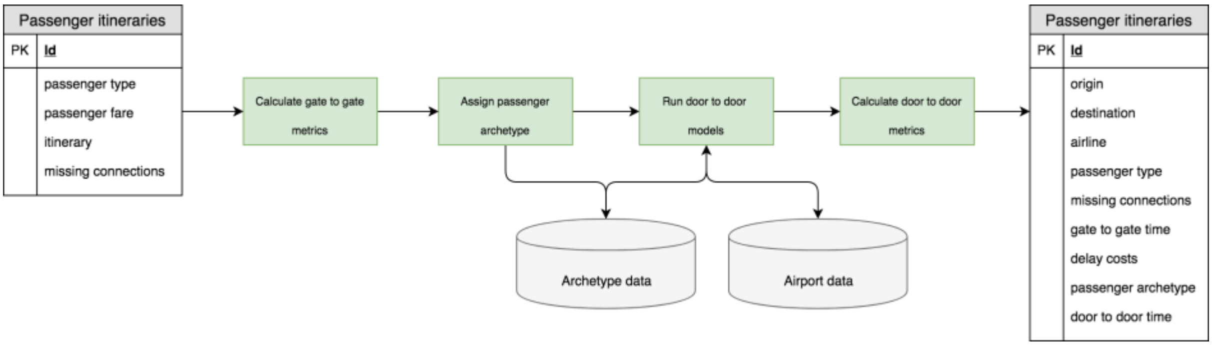

Figure 19. Processing pipeline from gate-to-gate to door-to-door itineraries

Firstly, the model calculates gate-to-gate metrics for each passenger itinerary, in this case the gate-to-gate time and the delay costs associated if the passenger experiences any type of delay. Next, the model assigns an archetype for each passenger type randomly. There are six different passenger archetypes in the door-to-door module and two type of tickets:

- **For the premium ticket passenger type:**
  - **Executives**: This group of passengers travel for business purposes, which requires them to reduce journey times, leading to only hand luggage being taken and using private cars or taxi as airport access modes.
  - **Exclusive experience travellers**: This passenger group travels for private reasons and is characterised by rather frequent travellers who are interested in short-stay trips and use public transport as a means of accessing the airport.
- **For the standard ticket passenger type:**
  - **Family and holiday travellers**: This particular profile consists of larger groups of more than three people, travelling for leisure purposes and staying away for more than 7 nights, which requires them to take several pieces of luggage.
  - **Best Agers**: Passengers within this category are mainly over 65 years old, thus usually retired and more flexible with regard to journey times than other groups, conducting short-stay trips with access to the airport by private car, either driving themselves or being dropped off by friends and relatives.
  - **Youngsters**: Travellers within this group travel for leisure purposes but are highly subject to budget constraints. This influences their choice of transport mode to the airport as well as the amount spent during the journey, cost savings are often more important than time savings. They use their technical devices as a point of contact with the airline and the services they need during the journey. They usually take cabin-size luggage to save money.
  - **Price-conscious business travellers**: Travellers within this group travel for business purposes but are subject to budget constraints. This influences their choice of transport mode to the airport as well as the amount spent during the journey, cost savings are often more important than time savings.

As soon as all passengers have an assigned archetype, passengers are grouped and then counted per origin/destination and archetype. This information defines the number of distributions that have to be modelled in order to assign door-to-door times for each passenger. This approach reduces the number of distributions pre-calculated and thus the performance is improved.

Door-to-kerb and kerb-to-door models assign a transport choice for each passenger following a generalised Bernoulli distribution of passenger transport mode preferences. Indeed, some transport modes may be discarded depending on each passenger archetype. There are nine different transport modes considered - each provided with a mean and standard deviation of travel time when such a mode is available at an airport. We then simulate the distribution of travel times depending on the type of transport. Each passenger is therefore randomly selected from the distribution of travel time depending on the mean of transport chosen that depends itself on the passenger.

$$T_m \sim \sum_{i=1}^{2} LogN(\mu, \sigma^2) \qquad T_m \sim LogN(\mu, \sigma^2)$$

Figure 20. Door-to-kerb stochastic models for different means of transport

Time of travel using buses and trains are modelled with a sum of two log-Normal distributions, which is the most common approach in the literature. The other means of transport such as car, car sharing or taxis are modelled as a (simple) log-Normal distribution.

The random selection also depends on the airport of origin or destination, respectively. Airports are categorised as:

- top ten;
- large seasonal;
- large non-seasonal;
- medium seasonal;
- medium non-seasonal.

Only the top ten airports are calibrated individually, as shown in the next section. There might be some transport modes that are more common than others for each categorisation. This will define which transport modal choice is selected and how long the journey will take.

The kerb-to-gate model generates a random time for each passenger. This random time depends on the following sequence of processes:

1. kerb walk;
2. luggage drop;
3. security;
4. immigration;
5. buffer.

The gate-to-kerb model generates a random time for each passenger. This random time depends on the following sequence of processes:

1. kerb walk;
2. baggage claim;
3. passport control.

Each process is calibrated individually using high-level information from the airports, specific studies or using the archetype calibrated mean and standard deviation of the processes. We have separated previously enumerated processes (except the buffer which has been modelled with a Normal distribution) into two types of processes, the queuing ones (e.g. luggage drop, security, etc.) that has been modelled using an exponential distribution and the others (e.g. kerb walk, baggage claim) that have been modelled with a truncated Normal distribution.

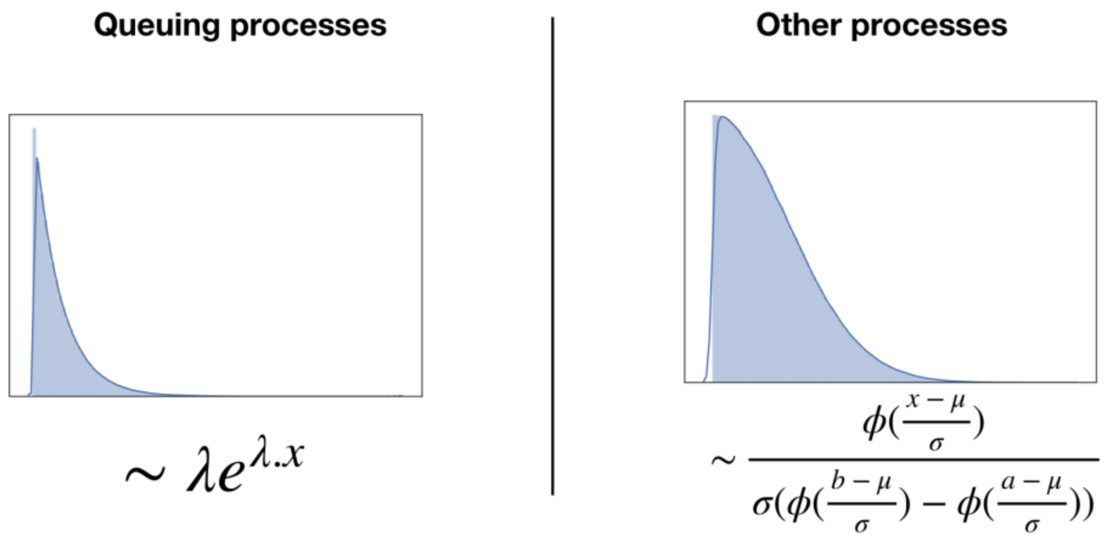

Figure 21. Stochastic models for gate-to-kerb and kerb-to-gate processes

Once each time distribution is created, the simulator assigns the door-to-door time for each passenger one-by-one, taking time samples for each distribution affected by the passenger. Note that this model assumes that passengers always arrive on time at the gate on their first itinerary leg.

3 Calibration

3.1 Strategic layer

3.1.1 Economic model

3.1.1.1 Calibration principles for the economic model

The calibration of the economic model is done on multiple sources of data which are far from consistent or aligned. We will highlight the main sources of data which have been used to set the different parameters of the model, in particular the initial state. The initial state of the model, is set to be as close as possible to 12SEP14.

3.1.1.2 Input data

(i) Airports

Regarding delay, data have been compiled from CODA data, with a breakdown of the different causes of delay, and DDR data. In the economic model we have focused on the turnaround delay (CODA codes: 85, 86, 87, 97, 99), which can be thought as the delay generated only by airports (in particular, congestion).

Following Section 2.1.1, we are interested in the relationship between traffic and average delay, which defines the capacity as one of the fitted parameters of the relationship. In order to fit this relationship, we needed more data than just an average per airport, as in the CODA data. As a consequence, we used some proxy, namely the actual off-block time minus the initial off-block time, as found in the DDR data (AIRAC 1409). The mean delay from the DDR dataset is then shifted to correspond to the one in the CODA data.

To do the regression, we first compute the proxy of delay on a given time window, for instance one minute, for the whole AIRAC. We compute the number of departures during the same time window and we perform a linear regression. We compute the goodness of the fit (using an R^2 score), and we then sweep the width of the time window and perform the regression for each of them. The width of the time window which fits best the data is selected and the corresponding parameters selected for this airport. Interestingly, the typical optimal time window is between 30 and 90 minutes, which can be regarded as the typical response time of the airport from a dynamical point of view. After this step, many airports have quite a high fit goodness, but some are still poorly fitted. We decided to set the parameters for them to the mean values of the best fits among the other airports. There might be a bias doing so, since the best fits are usually given by the busiest airports.

We also need a more detailed distribution for the delay in order to compute relevant expectations of the cost of delay. For reasons explained in Section 2, we decided to use a Normal distribution for the delays, even though a log-Normal distribution fits the data better. The Normal distribution is however a very good fit, with an average Kolmogorov-Smirnov distance of 0.09 compared to 0.06 for a log-Normal distribution. In terms of R^2, Normal distributions are around 0.98, whereas log-Normal distributions are around 0.99.

This fitting allows us to confirm that Normal distributions can be used in the model without too much loss of precision. Another problem is to estimate how the distribution changes with traffic. Assuming that the distribution remains Normal, and given that we already have the evolution of the average delay with traffic, it remains to address how the standard deviation changes. To do this, we conducted several analyses, comparing mainly the variation of the mean of the standard deviation. There are different ways of doing this:

- fit the distributions on different time windows and extract the standard deviations;
- directly compute the standard deviation from data on different time windows;
- calculate this for each airport separately;
- calculate this for all airports at the same time.

In the end, the best and most reliable results were obtained by computing directly the standard deviation on data for different time windows (2 hours). Moreover, instead of crossing the variations of the standard deviation with the traffic, we used the correlation with the mean. Behind this is the idea that there is some kind of homothety in the distribution with the traffic, which is the case in practice for many distributions arising from dynamical processes.

In any case, the most common way to find out if there is a relationship between two random variables is to use correlation coefficients. In other words, we consider the mean and variance as random variables linked to each other and to the traffic. These means and variances will then be used to build a distribution of delay for another random variable which is the delay.

Performing a correlation analysis, we find that for most airports, the correlation is positive, i.e. that the variance increases with the mean. Figure 22 shows the distributions of Pearson and Spearman coefficients among airports. Interestingly, a few airports exhibit a negative correlation, but looking more in detail at the data revealed that they are all very small airports, for which it is not clear why there is such a relationship. Note that in the figure we only plot the statistically significant coefficients. We also found that higher traffic airports have more correlated coefficients too, probably for some hidden reasons behind the distributions' behaviours.

Given this, we moved towards extracting some useful relationships for the model. As with the average delay, we used a linear equation to fit the standard deviation as a function of the mean (we also removed outliers using Mahalanobis distances). We find that in general the average R^2 is quite low, around 0.15. Since high-level traffic translates into higher correlations between mean and variance, we find that this score increases a lot for the very busy airports, between 0.70 and 0.95. We thus used the average relationship from these busy airports for those which have low correlations.

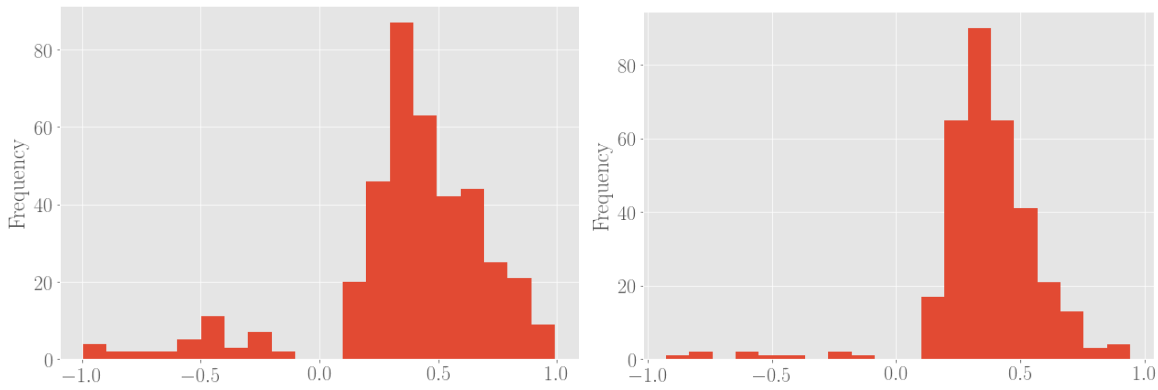

Figure 22. Distribution of correlation coefficients (Pearson, left, Spearman, right)

Distribution of correlation coefficients between mean and standard deviation delay across airports.

Cost and revenues for airports are then calibrated using data acquired from ACI EUROPE, including total costs, total aeronautical revenues, total non-aeronautical revenues for some European airports. Regarding cost, we decided to follow [3] and consider that a linear law is the best estimate without having access to better financial data for airports. Since operating the airport as a whole is somehow synonymous to delivering the capacity, we computed the linear term (c_1) by dividing the total costs of the airports by the capacity computed above. Regarding revenues, we computed:

- average aeronautical revenue per aircraft, which is used to compute the airline airport cost (see parameters P_O and P_D from Section 2.1.1).
- total revenues per passenger (which includes the average aeronautical revenue per aircraft), used to compute the revenues of the airports.

Once again, airports for which we do not have any data have their coefficients set to the average of those which do (the same possible bias applying, since we have financial data only for the main airports).

The final piece of information missing for airports is the minimum turnaround time, which is in fact used in the schedule mapper. We set it to 30 minutes uniformly for now, but it is planned to set them to better values, either from data analysis or another source of data.

Finally, common airport data are taken from various sources, including coordinates, time zone, etc.

(ii) Itineraries and legs

Initial itineraries are used to set up the model by fixing the initial volume and fares, as well as the possible connecting flights. These data have been produced and refined during other projects and comprise, in particular, the volume of two types of passengers (premium/standard) going from their original departure airport to their final destination, with all the intermediate connection. The average fare for each type of passengers is also available. We use this data to set the possible itineraries in the pax agents, and set the initial volumes and fares for each possible option.

The same data are used to set the individual legs that the Flight agent can operate, computing the average fare on the leg by considering only direct itineraries, and aggregating the different itineraries using this leg to have the initial volume of seats.

(iii) Passenger economic data

As highlighted before, the passengers are sensitive to different parameters, including price, income, and frequency of flight. In order to get the demand when one of these parameters change, we need some elasticities for each of them. The main reason why we chose a linear law is the fact that it is a standard assumption in economy and thus and that there is enough empirical data from the literature on the subject. For the price elasticity, we used the IATA economics briefing from 2008 [4] which gives different empirical elasticities at different levels of aggregation and for routes between different regions of the world. For income elasticity, we used the same report, also providing the income elasticities between different regions. We mapped each airport to its corresponding region and set the elasticities this way. For frequency elasticity, the literature is quite scarce, probably because it is more difficult to measure in reality. We found a reference [5] that report two elasticities, one for business passengers and the other one for leisure. We used these values overall in the simulations.

To setup the pax agents, we also needed their initial income, which is particularly important in the long-term scenarios. To do this, we downloaded a database from the World Wealth & Income Database (WID), which is freely usable for any purpose. Nearly all countries were represented in the database (in fact, all except Vietnam), so we were able to compute the average income for each country. A delicate matter was to assign this income to the passengers, because we did not know if they were leaving their own country or coming back. Since we know that richer countries have more travelling passengers (see for instance [6]), we used a linear rule of thumb to compute their respective proportions:

$$\frac{\mathrm{N_O}}{\mathrm{N_O}+\mathrm{N_D}}=\frac{\mathrm{i_O}}{\mathrm{i_O}+\mathrm{i_D}};\quad \frac{\mathrm{N_D}}{\mathrm{N_O}+\mathrm{N_D}}=\frac{\mathrm{i_D}}{\mathrm{i_O}+\mathrm{i_D}},$$

Where N_O, N_D are the numbers of passengers from the origin and destination country travelling between O and D and i_O, i_D are the corresponding average income. As a consequence, the average income on an itinerary from O to D is set to:

$$\mathrm{i_{OD}}=\frac{\mathrm{i_O}}{\mathrm{i_O}+\mathrm{i_D}}\mathrm{i_O}+\frac{\mathrm{i_D}}{\mathrm{i_O}+\mathrm{i_D}}\mathrm{i_D}$$

(iv) Airlines

For airlines, we first needed data on alliances. Since passengers can easily book a ticket with different airlines within the same alliance, it is very important for network effects to include them in the model. In the data we accessed (from in-house sources), airlines had different levels of collaboration: alliances, partnership, bilateral agreements, etc. For the need of the model, we used the following procedure:

- if two airlines belong to a partnership (or cooperation agreement), but only one of them belongs to a given alliance, the other airline is also assigned to the same alliance;

- all airlines in the same partnership without any other link to an alliance forms a new 'alliance';
- any remaining airline (without partnership or alliance) is part of its own new alliance (to ensure internal consistency in the model; it does not play any role).

The second step for airlines is to have information on their cost. In particular, we needed to compute the cost of a given flight plan flown by a given airline. In order to do this, we performed a quadratic fit of the fuel consumption as a function of the distance flown using BADA 4.1 model for different aircraft. Using the airline aircraft mix inferred from DDR data for each OD pair, we then computed the average fuel consumption coefficients for each of them, which were directly fed the Flight agent.

From BADA model, we also extract an average MTOW of each OD pair (in fact, an average $\sqrt{\mathrm{MTOW}}$). This was done to estimate the ATC cost, since the MTOW enters the formula, see Section 2.1.1. From the itineraries, we extracted the potential hubs of the airline. As seen in Section 2.1.1, the Airline agent needs a hub when it opens of new route if this route would have a connection (because the connection has to be made through this hub). In order to do this, we simply computed the airports through which at least one passenger makes a connection, with the same airline inbound and outbound.

We calibrated the cost of delay by using the coefficients taken from [7]. These coefficients are given per type of aircraft and per type of airline.

(v) Flight plans

Flight plans are used by the airline agent to compute the ATC cost and the fuel consumption. Flight plans were extracted from DDR data, analysing them so to have the distances flown within each ANSP.

(vi) ANSP data

The calibration process for ANSPs is slightly different with respect to the version presented in D5.1. In particular, the efficiencies of the ANSPs are set by internal calibration and not taken from the external source [8], since there are no other parameters to calibrate.

To calibrate the ANSPs, we used the assumptions that the ANSPs were at zero profit in the initial state. Using the previous flight plans and ANSPs' unit rates in 2014, we were able to compute their revenues. Using the fact that their capacities represent the maximum number of flight-hours that they can sustain, we calibrate their average cost per flight hours controlled by equating their costs and their revenues.

More specifically, for a given ANSP, we use their unit rate u and the historical traffic T to compute their revenues r, that we assume is equal to their cost c. Given a target delay, their capacity C is given by the inverse relationship (through numerical minimisation) as explained in Section 2. Their efficiency ω is thus c/C. We also set the maximum capacity to 120% of C, as also explained in Section 2, thanks to Belgocontrol's feedback.

Finally, we also flagged the non-ECAC ANSPs to be inactive during the simulation: they have a fixed unit rate throughout the simulation and we do not compute their costs, revenues, or profits.

3.1.1.3 Post-calibration checks

Post-calibration usually refers to the comparison of the output of the model with some empirical data in order to cross-check some facts and see whether there is any issue in the model. For instance, in Figure 23 we show a comparison of the cost structure of the airlines. On the left, we show the share of each kind of cost for two types of airlines (scheduled, low-cost) using empirical data [13], and on the right-hand side we show a similar plot obtained with the model, using the 'current' scenario.

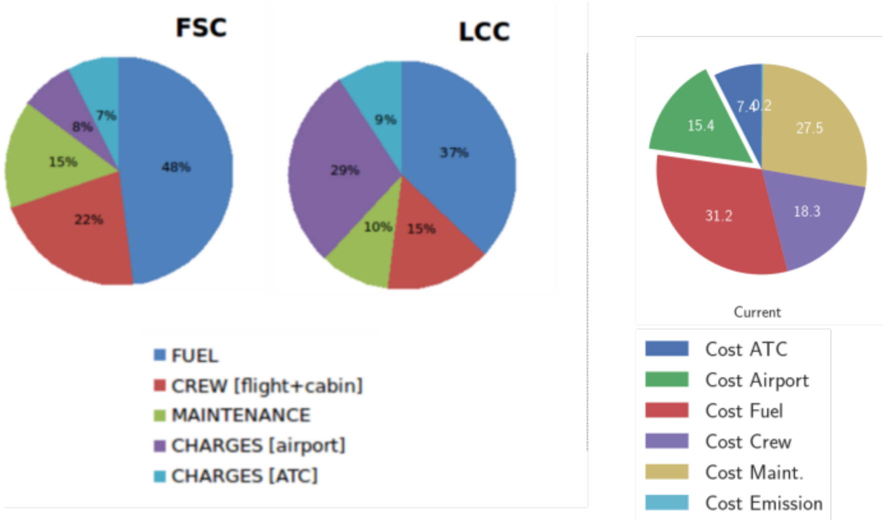

Figure 23. Airline cost data comparison between empirical data and outputs of the model

Empirical data source: ICAO DATA+ air carrier finances [13].

Keeping in mind that the right-hand side is an average over all the types of airline, the agreement between both is quite good. Some differences still exist, for instance the cost of fuel seems to be underestimated in the model. The ATC and crew costs seem to be robust, whereas the airport charges have too much variability on the left to compare them efficiently with the model. The maintenance costs seem to be potential overestimated in the model, although the inputs were validated with several airlines.

Once a model has been used to calibrate input data, there is usually a second phase where some free parameters are tuned so that the output is matched to empirical features in the data. This is also related to sensitivity analyses, where one changes one parameter (usually not calibrated on data) to see how it affects the model.

There are few free parameters in the model, most of the others being already set with the pre-data analysis. To properly explore the model, further heavy simulations would be needed and are planned for further research. The present results have been obtained with 'manual' calibration, exploring a small part of the parameter landscape and setting them in order to have a realistic output and a stable equilibrium.

3.1.2 Schedule mapper

We saw in Section 2.1.2 that the schedule mapper relies heavily on historical information. In input, the schedule mapper needs four types of information:

- airport data: standard airport data, also used in the economic model and elsewhere;
- historical schedules: the schedules that were already planned for 12SEP14;
- strategic input: the flows from the economic model;
- pattern and taxon data: the historical taxons and patterns, both defined in Section 2.1.2.

All data except the third category comes from a prior analysis of historical data. We will in particular focus on pattern and taxon analysis. All this analysis in based on DDR data, with 6 days from, AIRAC 1409. To have cleaner results, we used only flights classified 'S' (for scheduled) and considered only European airlines (otherwise an aircraft can be omitted because it links two non-European airports). This reduces the number of flights from around 118k to around 21k. However, this filtering process might be too selective, and will be refined further in the next version.

3.1.2.1 Taxons

First, we look at the open and closed patterns in the data. Note that all patterns on a sufficiently long time window should be closed (i.e. the aircraft always returns at a given airport at some point). Over 6 days, many tail numbers end up forming open patterns. However, most of these tail numbers are used only once. When counting the number of flights operated, most of them (72%) are part of at least one closed loop during the 6 days period, 39% when we reduce the time-window to 72 hours, as in the following. The distributions for 6 days are shown in Figure 24. This shows the number of tail numbers having at least one closed pattern, or none (left), and number of flights operated by an aircraft making at least one loop, or none (right).

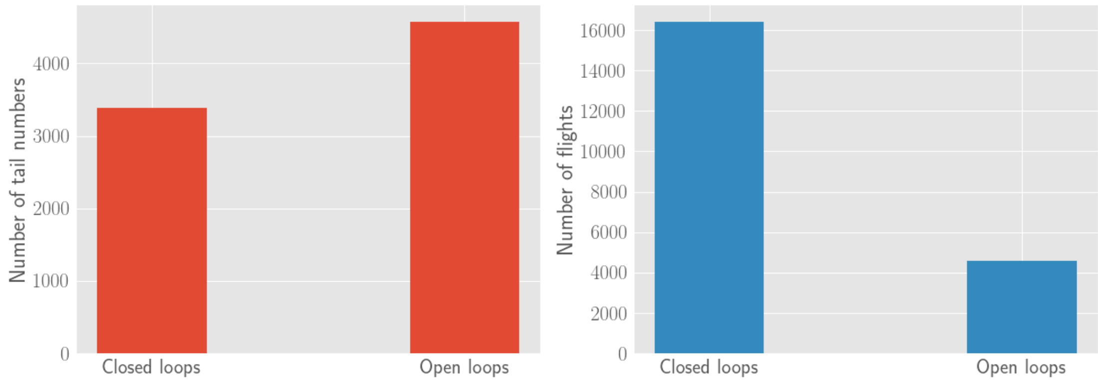

Figure 24. Number of tail numbers (left) and flights (right) by closed or open loops

Once we have all the closed loops, we can begin the taxonomy classification. As explained before, several patterns can represent the same taxon if their unlabelled succession of airports are similar. In fact, we even described the elementary taxons in terms of a succession of distinct airports. As a consequence, there is as many elementary taxons as the maximum number of airports in a loop (minus one), see Figure 3 for the first three. Figure 25 shows how many of these elementary taxons can be found in the data.

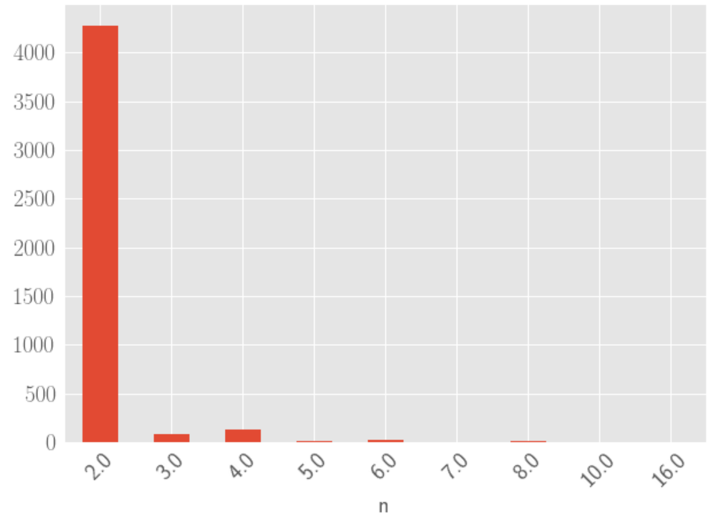

Figure 25. Distribution of elementary taxons, characterised by number of airports (abscissa)

The vast majority of these taxons are represented by the first type, i.e. the back and forth, a few others using the second or third taxon, see Figure 3. As a result, in this analysis we decided to focus only on the elementary taxons with a maximum of 4 airports, which capture more than 99% of them. As a consequence, we can consider these three taxons as the building blocks of the airlines' networks, and in the model we made a further simplification by considering only the taxons composed by the first elementary taxon ('back and forth'). As explained before, these 3 elementary taxons can then generate a high number of taxons by combining them. It is thus important to see what the most represented taxons are in the data, which is represented in Figure 26.

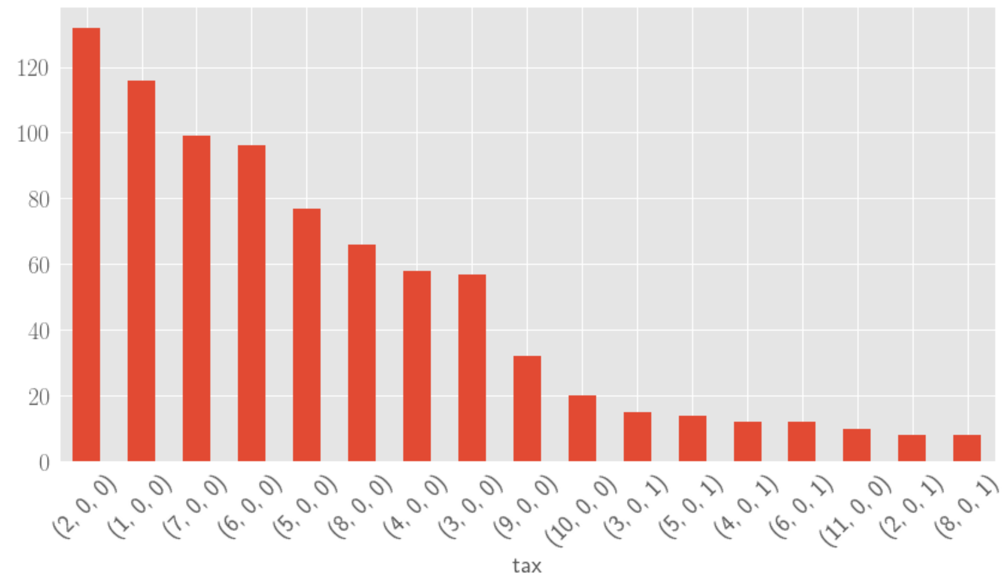

Figure 26. Distributions of the most important (17 first) combined taxons

Once again, many taxons are poorly represented (not in figure), but there is however a rich taxonomy, with 17 taxons needed to account for 90% of the flights. The two most common are (1, 0, 0) (one back and forth in 72h) and (2, 0, 0) (two back and forths). Almost all of them represent several back and forths between the same two airports, with a few of them using a quadrangular pattern (0, 0, 1). Interestingly, the triangular pattern is seldom used, its first appearance being in 18th position (outside this figure).

3.1.2.2 Departure and turnaround times pre-analysis

Another issue is how to take into account realistic turnaround and departure times. In practice, these are defined by crew and maintenance, as well as airport slots. Since we do not have enough information on these, we try to take them into account by analysing the empirical turnaround and departure times. The following is an illustration of the initial analysis, but it led to a different use in the final version of the model (see below).

When doing so, it is important to keep different distributions for the different taxons. Indeed, 3 back and forth in 72h represents a daily flight for instance, and so the turnaround time would be picked around the duration of the night. A single back and forth route in 72h is probably more due to transatlantic flights for instance, and might have different turnaround times.

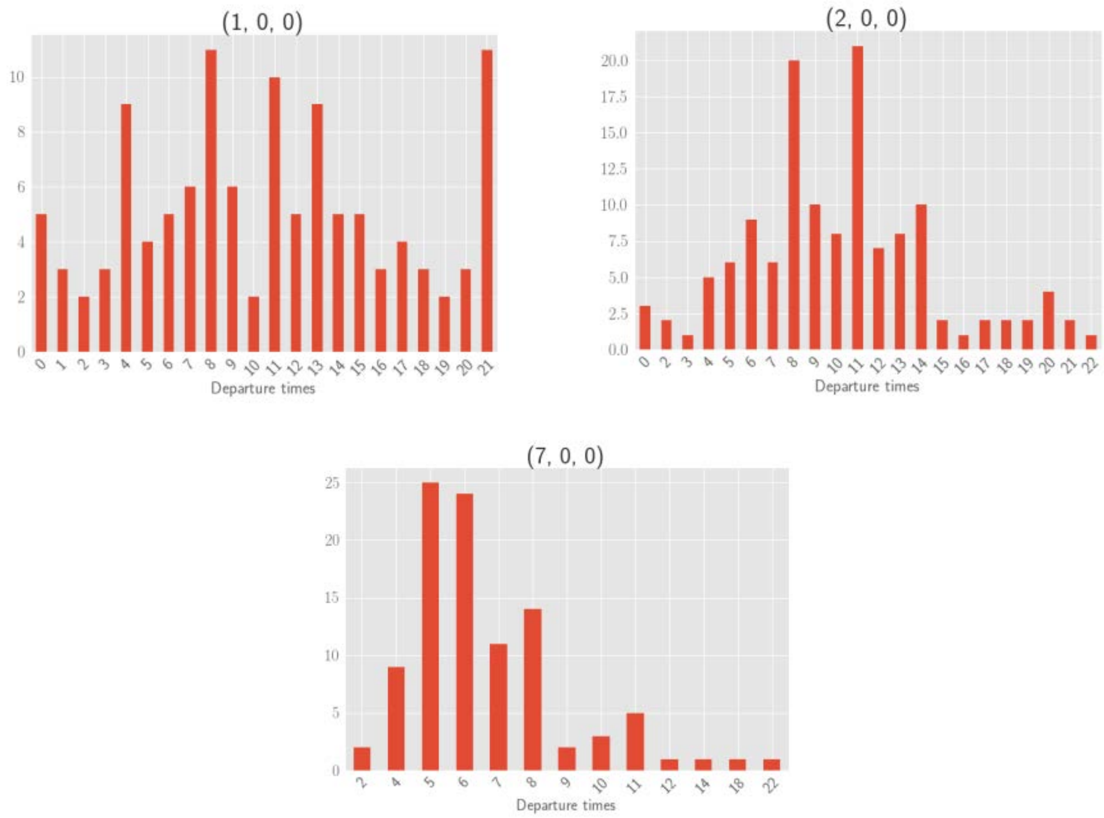

Figure 27. Departure times of aircraft (from first airport) for taxons (1, 0, 0), (2, 0, 0), and (7, 0, 0)

Figure 27 shows the difference in distributions for the three first taxons. The aircraft doing only 1 back and forth in 72 hours for instance have quite a broad distribution but have peculiar peaks at

0000, 0400, and 2100, whereas aircraft doing only two back and forths seem to depart earlier, around 0500 mostly. These peaks are completely absent from the distribution of taxon (2, 0, 0), where most of the flights depart in the morning. Taxon (7, 0, 0) has quite a well-defined distribution, with most of the flights departing in the very early morning.

Connecting times are equally important and are also collected and grouped by taxon type. Figure 28 shows their distributions for the same three taxons as previously. Taxon (1, 0, 0) has quite high turnaround times, with a significant proportion of them greater than 10 hours. This is expected, since these flights would be usually very long-haul and thus have more maintenance time and more complicated schedule constraints. Taxon (2, 0, 0) has much smaller turnaround times, with a few peaks after 10 hours, corresponding most probably to maintenance times. Finally, taxon (7, 0, 0) has quite a different distribution, with a main peak around 1 hour, corresponding to very fast turnaround times, and a smaller one around 8-9 hours, corresponding most probably to night turnaround times.

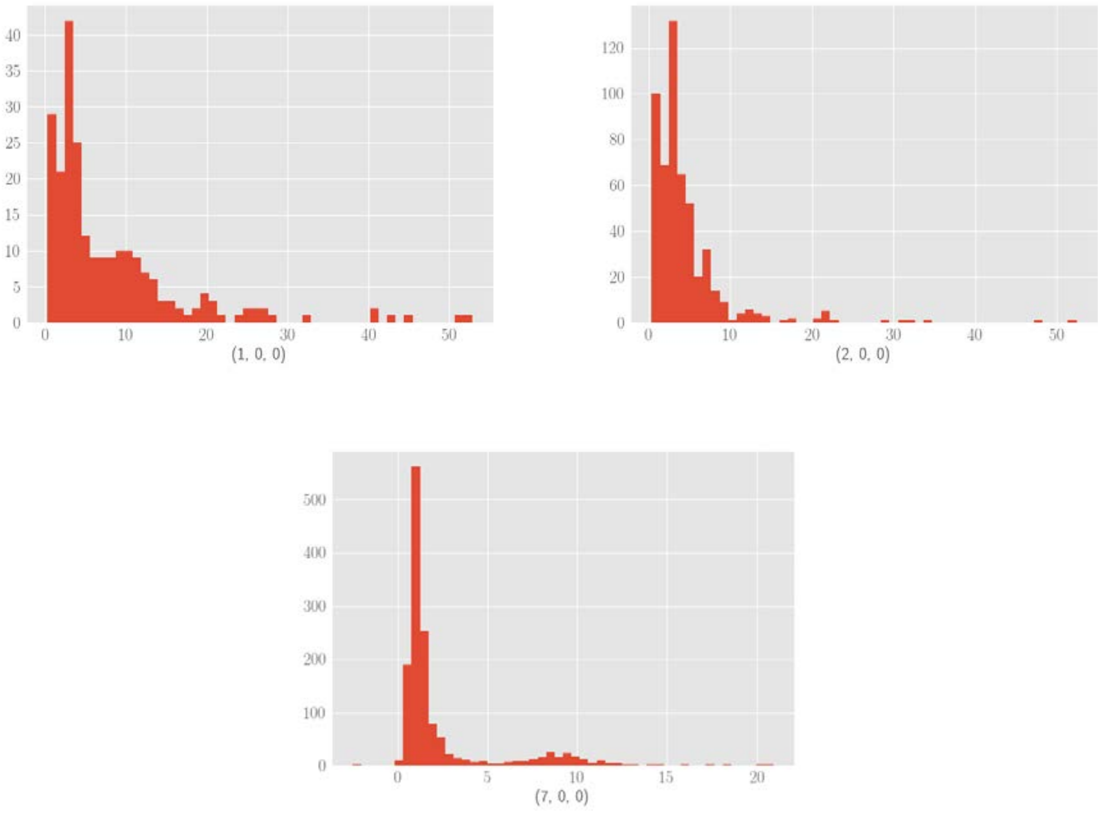

Figure 28. Distribution of turnaround times (in hours) for taxons (1, 0, 0), (2, 0, 0), and (7, 0, 0)

3.1.2.3 Decision tree for turnarounds

In the final version of the model, the schedule mapper uses only one type of elementary taxon and more parameters need to be taken into account to choose turnaround times, compared with the earlier version presented in D5.1. The schedule mapper does not choose the taxon before the initial departure time; we draw at random a departure time when the first edge is chosen (see Section 2). The type of pattern then emerges naturally from the departure time, the distances to travel, etc.

Regarding turnaround and departure times we use a decision tree. Decision trees are powerful machine learning techniques where one tries to recursively find the best parameters explaining the variance in the data. It is able to seamlessly include categorical and numerical data. It is also a white box, i.e. it can be transparently explored, even by non-experts.

We select historical data containing departure times and associated data, such as the type of airline. Then we select part of these data for training (75%) and the other part for validation (25%). We then choose to use an R^2 score (using the difference between the predicted and the real turnaround times) on each subset to estimate the goodness of interpolation and extrapolation, respectively.

We then need to choose the predictors, i.e. the parameters whose values will be used to predict the departure times. As a general rule in machine learning, the more predictors are included, the better is the goodness fit on the training data. However, having too many predictors leads to poor extrapolation power, as indicated by decreasing R^2 scores in the validation data. After manually trying different predictors, we found that the following set of variables gives the best results:

- the previous time of arrival (numerical data);
- whether the airport is in the ECAC space (categorical, two categories);
- which type of airline is operating the flight (categorical, five categories: scheduled, low-cost, charter, regional, and others);
- the number of passengers on the flight (numerical data);
- the distance flown during the previous flight (numerical data).

An illustration of the tree is shown in Figure 29 and Figure 30. It reads as follows. Starting from the most upper node, we see that the most important variable to explain the variance in the data is 'a_time', which corresponds to the previous arrival time. Since the data have been standardised, the threshold value is not (human) readable, but the conclusion is that if the time of arrival is below this threshold, we need to continue by going through the left branch to the next node, otherwise we need to hop onto the node on the right branch. Then we do the same for the other nodes. For instance, on the left, the next most important variable to explain the variance in this subset of data is the distance flown previously by the aircraft (also normalised). One can go down like this to the 'leaves', where the final variance and the number of samples are indicated.

Once this tree has been trained, we save it and the schedule mapper can use it. To do that, it only needs to input the above variables in the tree and get back the relevant leaf. It then chooses one of the turnaround times present in the samples of this leaf.

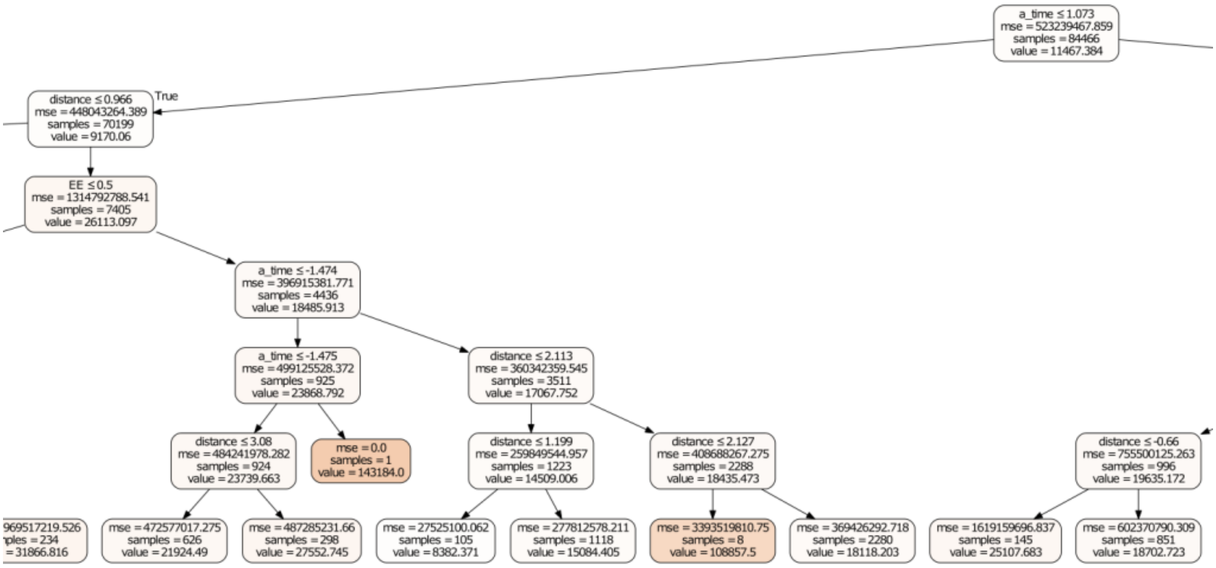

Figure 29. Upper part of the decision tree for turnaround times

For each node, if the criteria in the node is true, one needs to follow the left branch. Categorical data appear with numerical inequalities. For instance, EE <0.5 means 'not in ECAC space'.

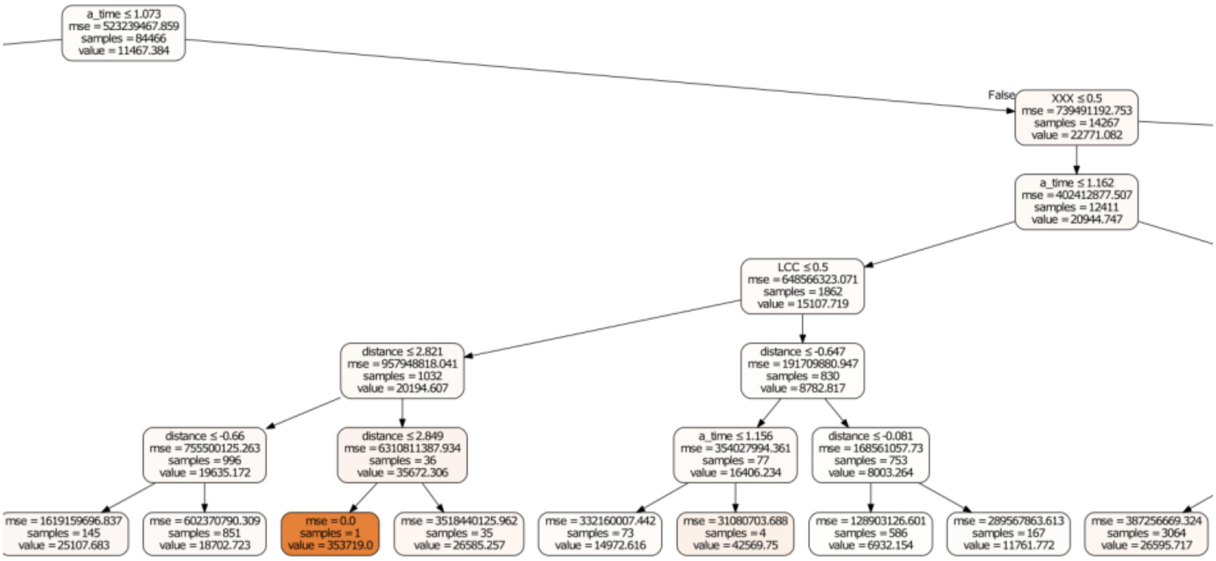

Figure 30. Lower part of the decision tree for turnaround times

This shows some instances of choice on the type of airline. For instance, 'XXX < 0.5' means that the airline should not be of type 'others'.

3.2 Pre-tactical layer

3.2.1 Flight plan generator

3.2.1.1 Calibration principles for the flight plan generator

As described in Figure 6, the flight plan generator relies on the analysis of historical data. This historical data analysis is used to generate a pool of flight routes between origins and destinations and to generate distributions that are used to model flight trajectories (climb, descent and cruise phases characteristics) and weather factors (cruise wind).

The data analysed is primary sourced from DDR (AIRACS 1313-1413, when detailed information is needed 5 individual DDR days from AIRAC 1409 are used) and BADA performance models.

3.2.1.2 Input data for flight plan generator

(i) Historical flight plan routes

The number of possible routes between origin and destination are computed based on historical last submitted flight plan data from the DDR data sets (m1). These trajectories include the ANSPs used and the distance flown within each ANSP. As expected, flights use the same or similar routes consistently between the same origins and destination. However, in some cases the data includes slightly variability which is not a reflection of a difference in routes but of the operations of the day. A clustering of historical routes is carried out to identify the distinct possible routes that will be included in the pool of historical routes for each origin and destination pair.

The clustering should capture all the possible distinct routes between each origin and destination airport which either use different NAS airspaces or are significantly different in length. For this reason, all the routes between the same origin and destination are first grouped by the used NAS in the order they are crossed. Then for each of these groups of routes a clustering is done using Kernel Density Estimation (KDE) based on the total route length and keeping the local maximum of the KDE as the representative of the class. This allows us to identify significant difference in length of routes that use the same airspaces at a NAS level.

For example, Figure 31 a) presents all the historical routes between EBBR and LTBA (in this case there are 26 flights in the period analysed). Figure 31 c) presents the result of the application of the clustering reducing the number of routes to 12. As presented in Figure 31 d), the KDE is applied for each of the routes using the same airspaces. For this reason, there are as many groups (arches in the plot) as the possible number of different sequences of NAS used. The figure represents for each group of NAS used the density distribution as a function of flight plan distance, the local maxima (in green) and minima (in red) are represented in the plots; local maxima are considered representative of the categories for clustering purposes. The different flight plans are marked as blue crosses.

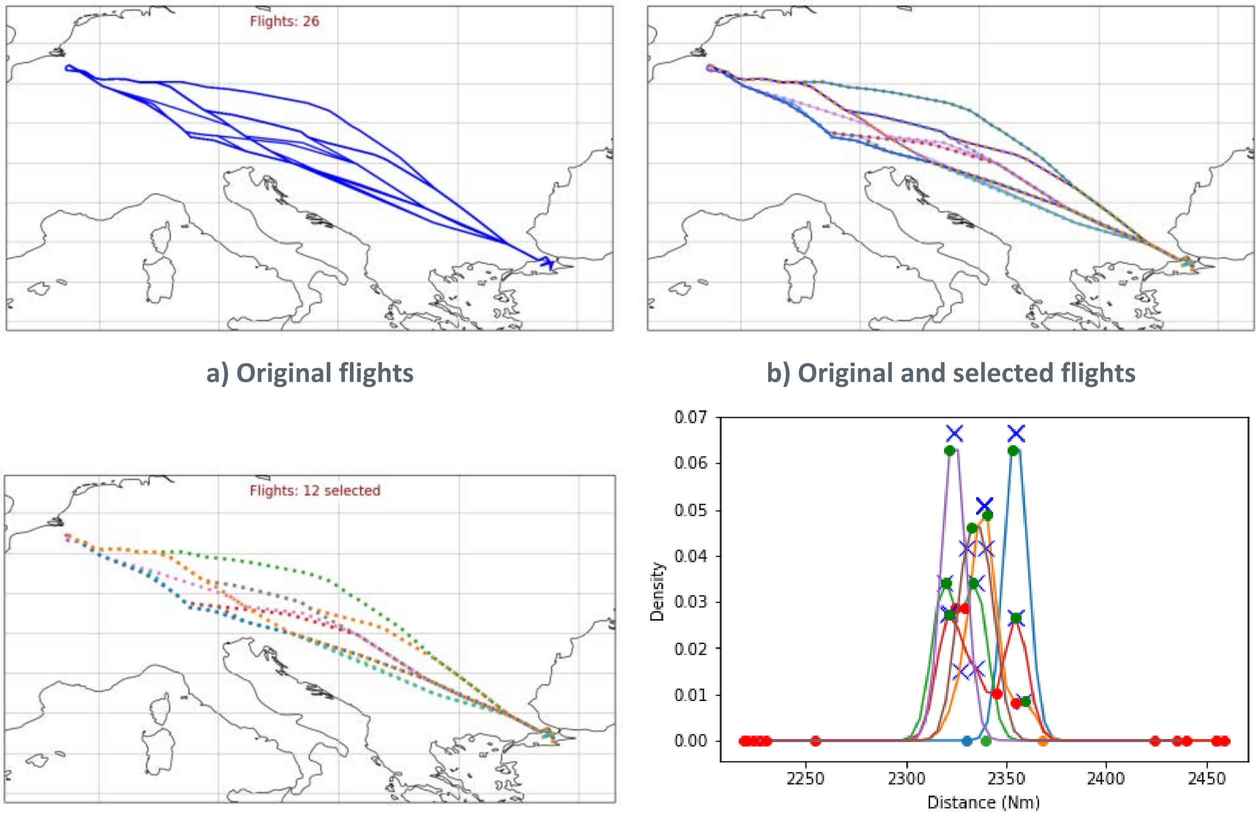

a) Original flights

b) Original and selected flights

c) Selected routes

d) KDE clustering

Figure 31. Example of route clustering EBBR – LTBA

Figure 32 presents another example between DGAA and EGLL. In this case, t is easier to observe in Figure 32 c) how there are two possible NAS usage (DG-DR-DA-LF-EG and DG-DR-DA-LE-LF-EG). In the DG-DR-DA-LE-LF-EG case two possibilities are kept there (a route with at total length of 2903 NM and a route with a length of 2867 NM) while in the DG-DR-DA-LF-EG only one route is maintained (with a length of 2901 NM). Some of the flights using the same airspace do not differ in distance enough to consider them separately and in this manner the original 10 flights between DGAA and EGLL are reduced to three possible routes.

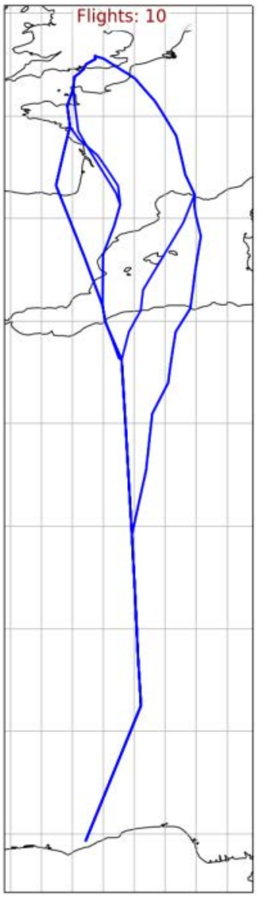

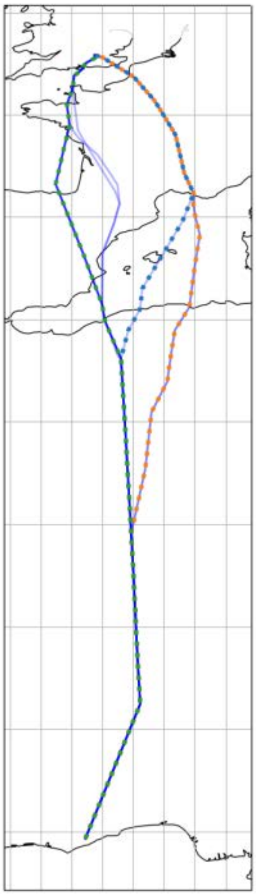

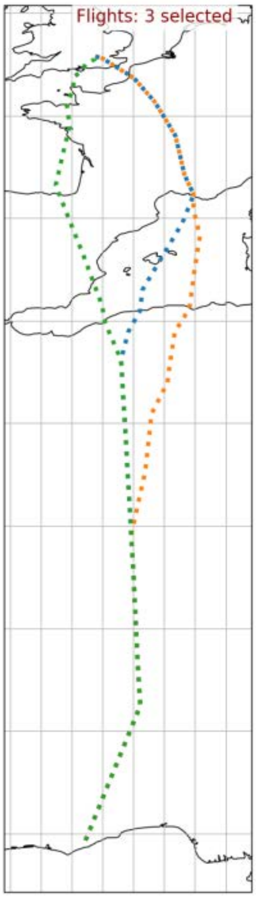

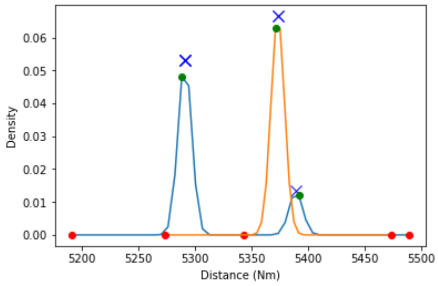

a) Original flights b) Original and selected flights c) Selected routes d) KDE clustering

Figure 32. Example of route clustering DGAA – EGLL

(ii) Historical flight plan requests

The DDR dataset contains information on the flight level and speed requested by the flights. A given flight might request different flight levels and speeds for different parts of its trajectory. Figure 33 shows two examples of trajectories with their associated speed an FL requests.

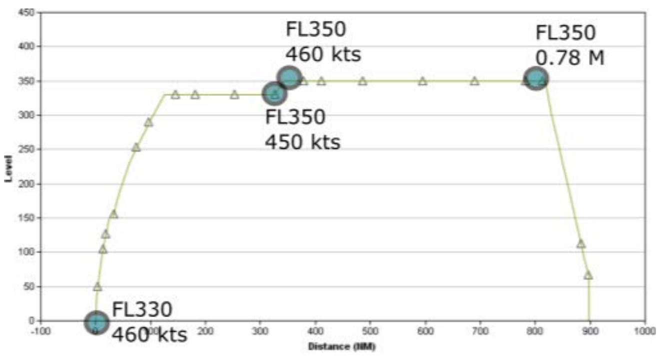

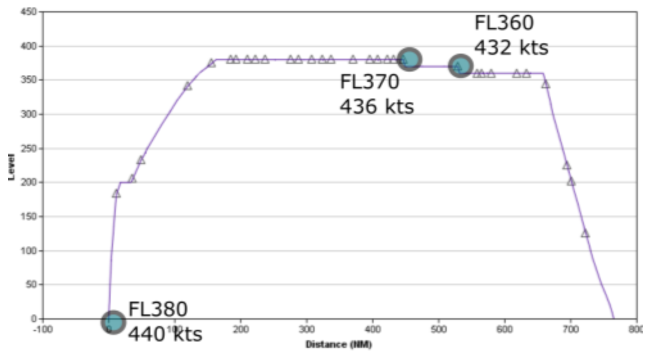

a) LEMD – GCRR b) LIEO – EDDT

Figure 33. Example of flight level, speed request

An estimation of the average FL used is computed by averaging the FL requests weighted by the distance for which that flight level will be used. Figure 34 shows the average flight level computed for the examples of Figure 33.

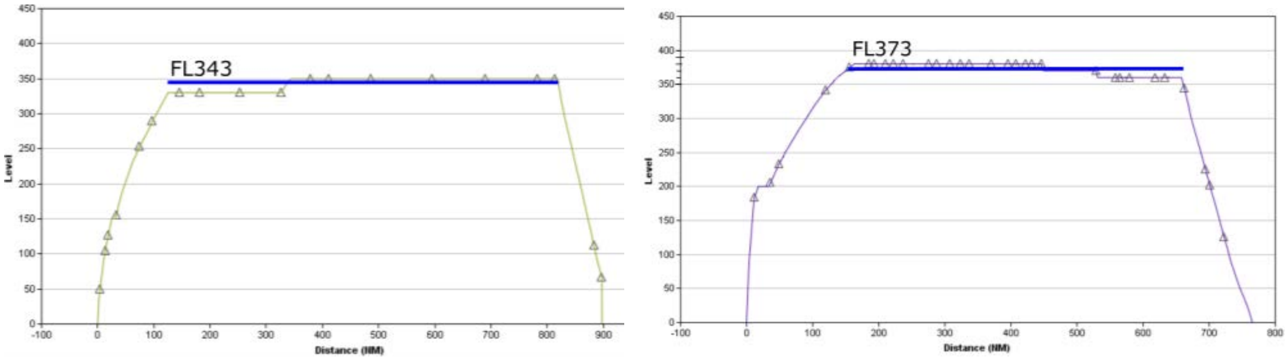

a) LEMD – GCRR b) LIEO – EDDT

Figure 34. Example average cruise flight level

The flight level requests are then classified by aircraft type and by flight plan distance category (very short, short, medium, long, very long, extra long) to generate cumulative probability functions as shown in Figure 35 for the B737.

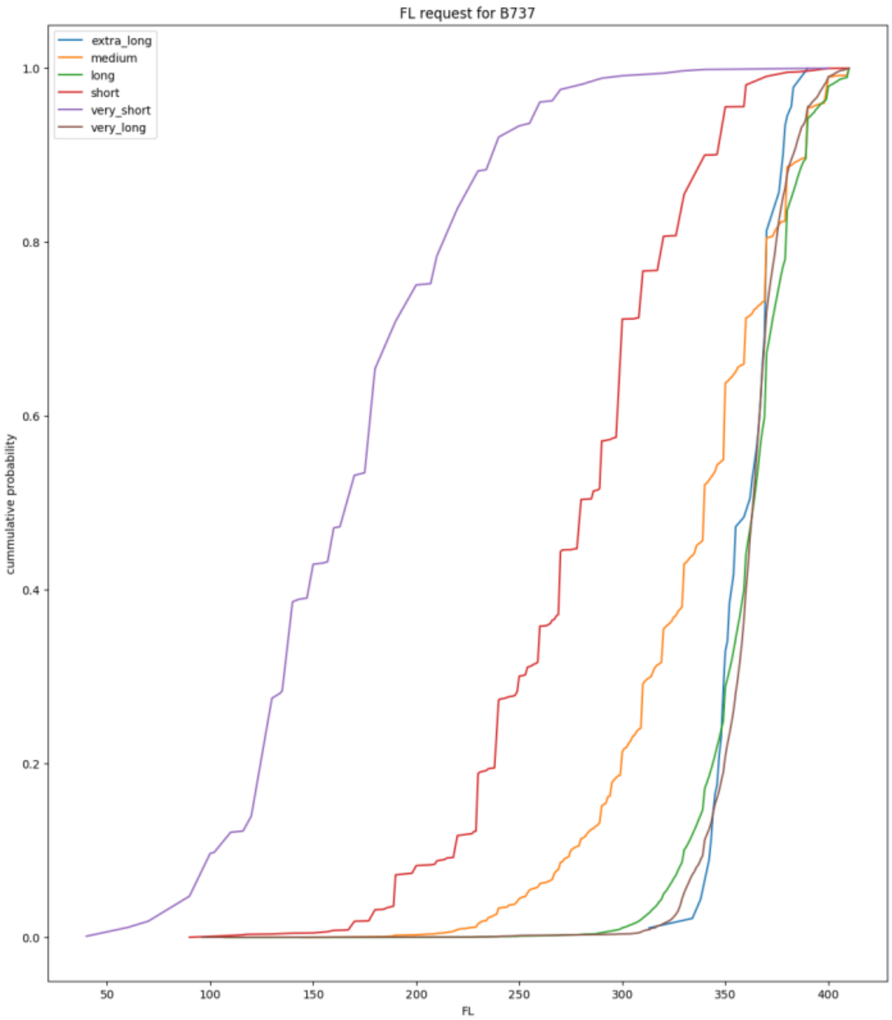

Figure 35. Example distributions of average cruise FL requested for B737 by FP distance

The speed requested are in turn grouped by the flight level they have requested to give us a distribution per aircraft type and flight level as shown in Figure 36 for the A320.

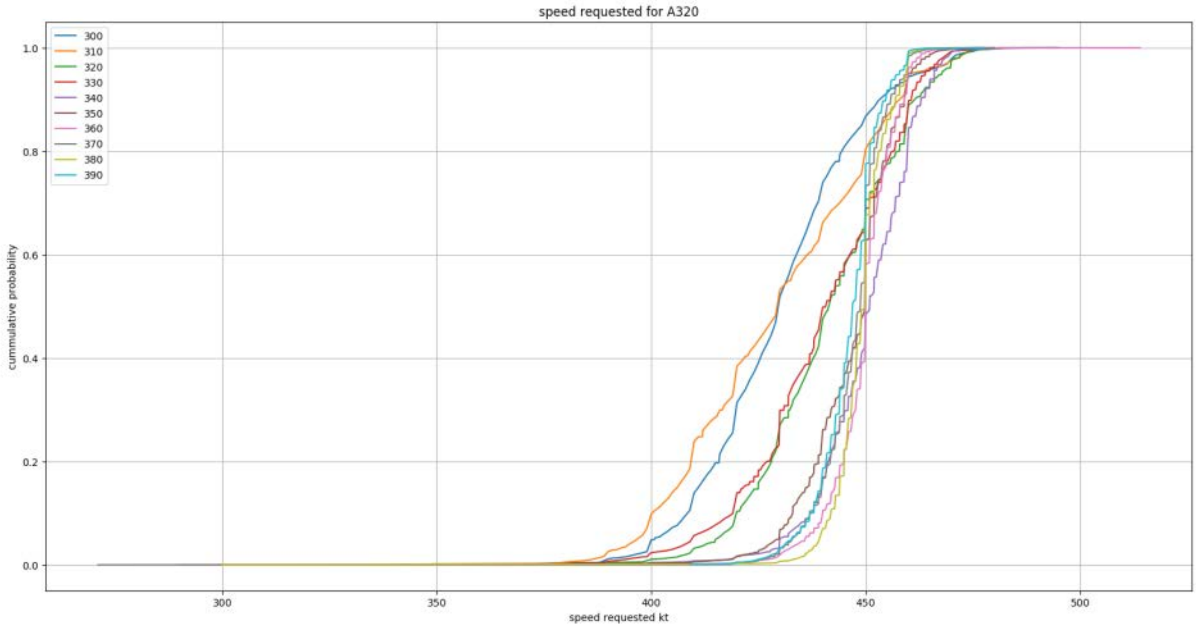

Figure 36. Example of distributions of speed requested for A320 by flight level

(iii) Historical flight plan trajectories

The full trajectory contained in the DDR dataset is analysed in its vertical profile to obtain distributions of distance and time required for the climb and descent phases.

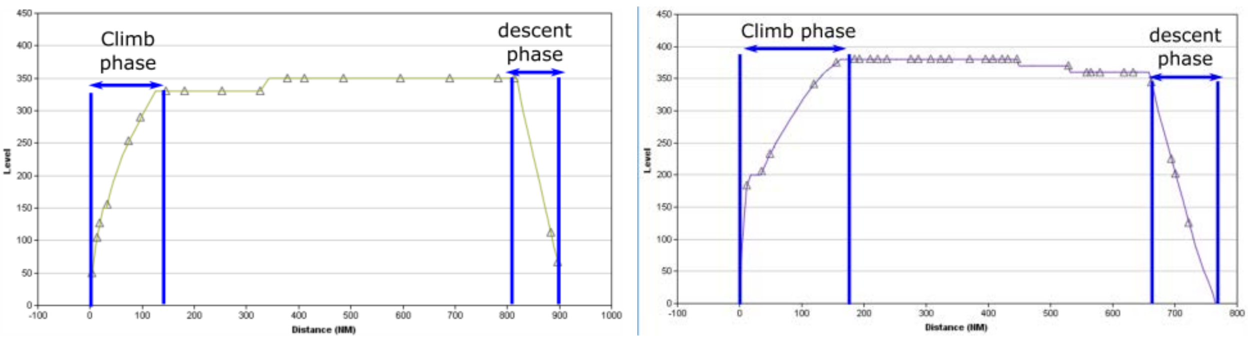

a) LEMD – GCRR b) LIEO – EDDT

Figure 37. Example of estimated climb and descent phase

Firstly, the end of the climb and the beginning of the descent are identified in the trajectories. This is done by considering the passage of the first flight level requested for the climb and the descent from the last flight level request. Figure 37 shows the climb and descent estimated for the two example flights.

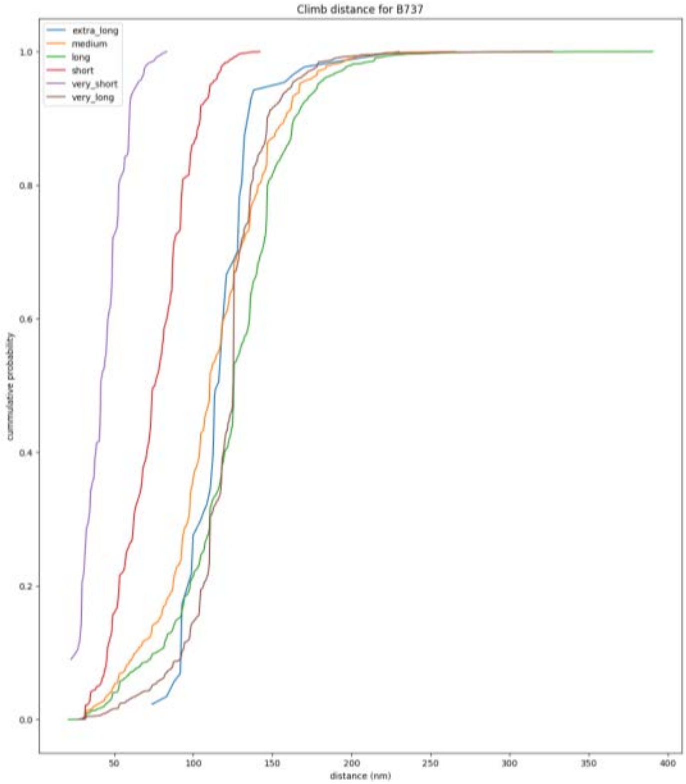

a) Climb distance

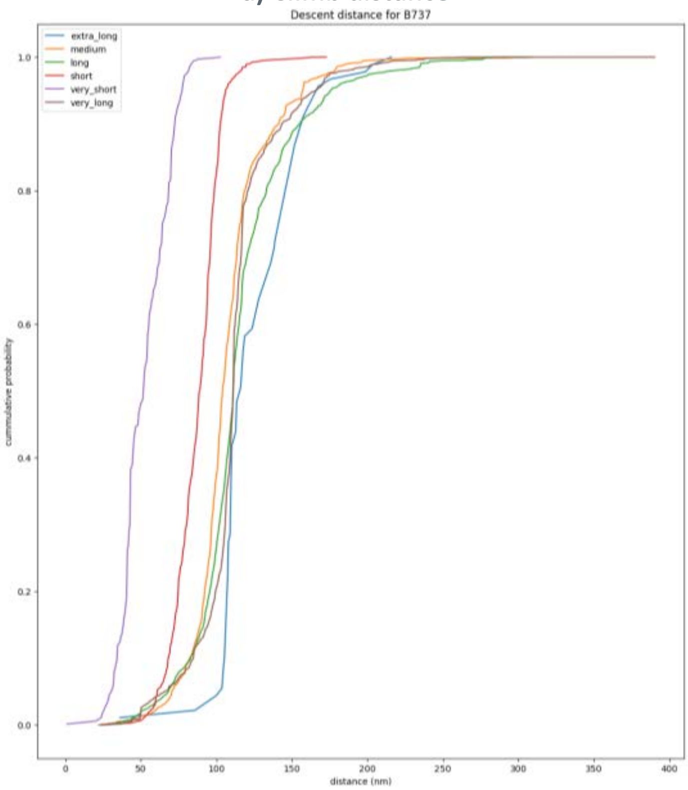

b) Descent distance

Figure 38. Example of distributions of climb and descent distance for B737 by FP distance

These climb and descent distances are grouped in a similar manner as the flight level requests to generate cumulative probabilities functions per aircraft type and flight plan distance category. Figure 38 shows these probabilities for the climb and descent distances for the B737.

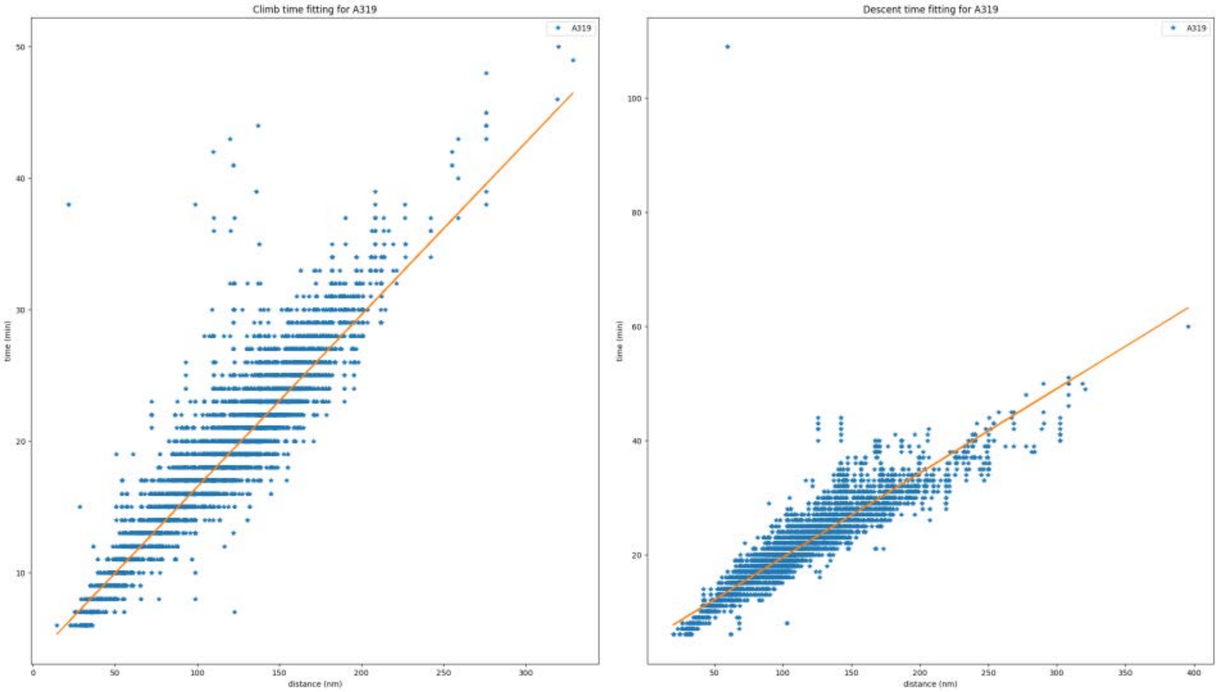

a) Climb time fitting (R^2=0.87) **b) Descent time fitting (R^2=0.83)**

Figure 39. Example of linear fitting of climb and times for A319

For each aircraft type a linear fitting between the climb and descent distance and the time required to perform them is computed. Figure 39 shows this fitting for the A319 aircraft. Note that even if the R^2 is high there is still some variability. For this reason, the standard deviation between the fitting and the actual data is used in the model to add some variability to the times computed with the linear fitting, as explained in Section 2.2.1.2. Note that aircraft types have been clustered based on their performances to increase the data per aircraft.

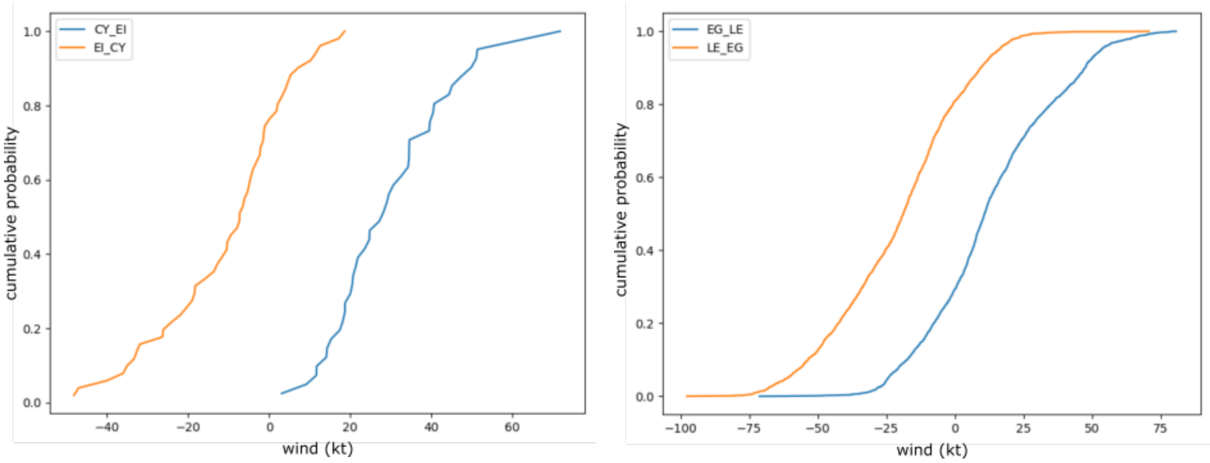

a) Flights to-from Canada–Ireland **b) Flights to-from United Kingdom–Spain**

Figure 40. Examples of average wind distributions between NAS

It is possible to identify segments in the historical trajectories which are level at a requested FL. An average estimated wind can be then computed as the difference between the estimated ground speed (distance of the segment divided by time required to cover the segment according to the flight plan) and the requested speed. A weighted average wind is then estimated per flight considering the length of all the segments. Finally, cumulative distribution functions of average wind are produced by grouping the flights based on their origin-destination NAS. As shown in Figure 40, we obtain, in this manner, probability distributions of average cruise winds encountered, which differ by origin and destination, thus capturing the general weather pattern (e.g., flights from Canada to Ireland will in general have tail wind (positive) while flying from Ireland to Canada head winds are more commonly encountered (negative winds), see Figure 40 a).

(iv) Aircraft performance models

The pre-tactical block uses BADA performance models to estimate fuel consumption during the different flight phases and the different aircraft weights. When possible BADA 4 performances are used, if the aircraft model is not available in the BADA 4 dataset then BADA 3 is used.

3.2.1.3 Post-calibration of flight plan generator

The objective of the post-calibration is to analyse and ensure that the flight plans that are generated are as close as possible to historical flight plans when the scenario is set to the current scenario. Ensuring that the results are close to current operations gives us confidence on the model results.

For the flight plan generator there are different metrics that can be analysed: flight plan distance, buffers, fuel usage and en-route costs, route preference between origin and destinations and demand at the different NAS. We can calibrate the model to ensure that these metrics are as close to historical values as possible obtaining in this manner routes that are realistic.

The schedules that have been used in the baseline scenario are those from the flights of 12SEP14. The flight trajectory generator and the flight plan generator have been executed once then for all the flight plan options, the Flight Plan Selector has been executed a total of 50 times. The fuel cost used is a base cost of 0.5 EUR/kg of fuel and the CRCO charges are as set in FEB18 [12].

(i) Flight plan length

For each origin and destination different routes are possible, the Flight Plan Selector selects the flight plan to be used per schedule. Figure 41 shows the comparison of the distance of historical flight plans compared with the average distance of the flight plans selected per schedule. As observed values are very similar and differ from the distance of all the flight plan generated (Figure 42).

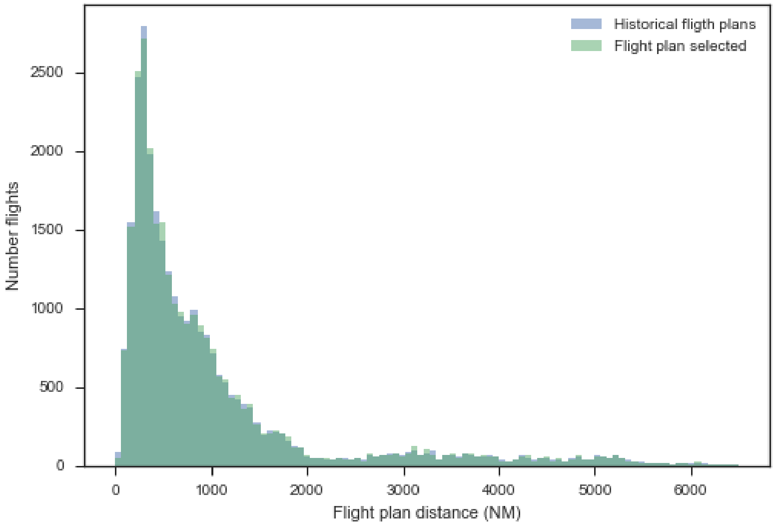

Figure 41. Flight plan distances: historical flights and average for FP selected per schedule

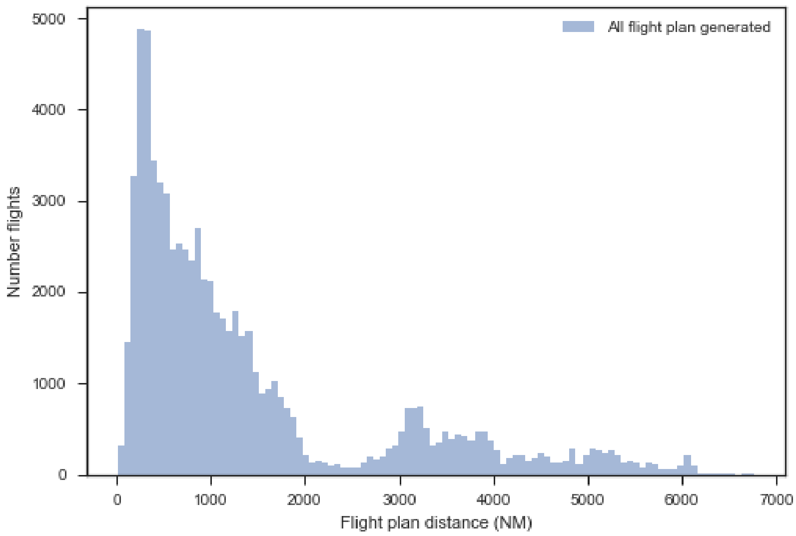

Figure 42. Histogram of flight plan distance for all FPs generated

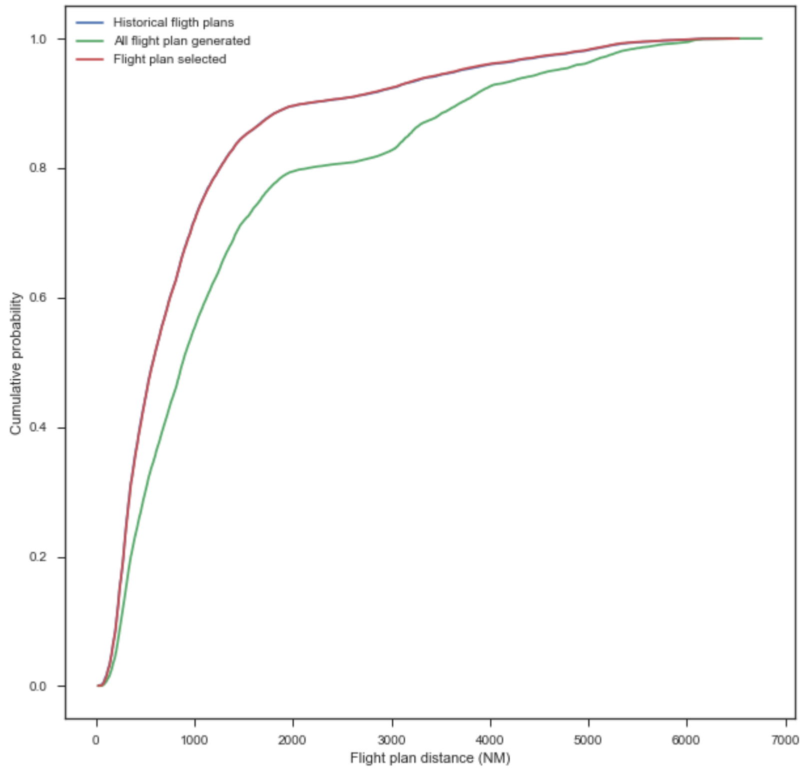

Figure 43. Cumulative probabilities of FP distance (NM)

Figure 43 shows the cumulative probability of having a given flight plan distance. As observed, all the generated flight plans have a larger flight plan distance than the historical and the selected flight plans. This shows that the route generator and the trajectory generator are creating options that are longer than those that are finally selected. This reflects also the fact that for longer flights there are more possibilities and hence there is a higher number of observed longer possible flights. There are more possibilities for longer flights but the flights selected are similar in length to historical flight plans. Table 4 shows some statistics on the flight plan distance for the historical flights, all the generated and the selected.

Table 4. Quantiles and mean for FP distance (NM)

Type	Quantile						Mean
	0.1	0.2	0.5	0.7	0.8	0.9	
All flight plans generated	255	354	882	1423	2160	3695	1407
Flight plans selected	208	280	579	958	1259	2135	960
Historical flight plans	207	279	557	956	1256	2147	961

(ii) Buffers and planned taxi times

As shown in Figure 12, for each flight schedule we have a flight block schedule defined as SIBT - SOBT, and if we compare that value with the flight plan time from take-off to landing we obtain the planned buffer time allocated to taxi-in, taxi-out and schedule buffer to not generate delay with that flight. Note that these buffers are planned and dependent on the flight plan selected. In some cases, tactically some delay can be recovered so a negative buffer does not imply directly an arrival delay in the flight but a planned arrival delay if the flight is executed as in the flight plan.

The schedule buffers are hence dependent on the flight route selected but also on the plan generation including which cruising speed is selected and the winds aloft.

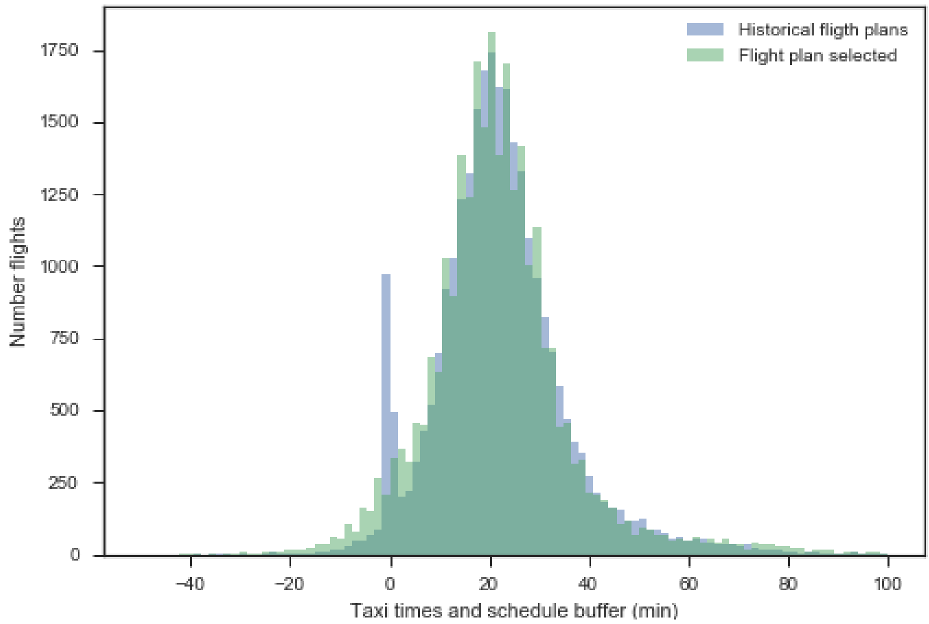

Figure 44. Taxi times and schedule buffers for historical and selected FPs

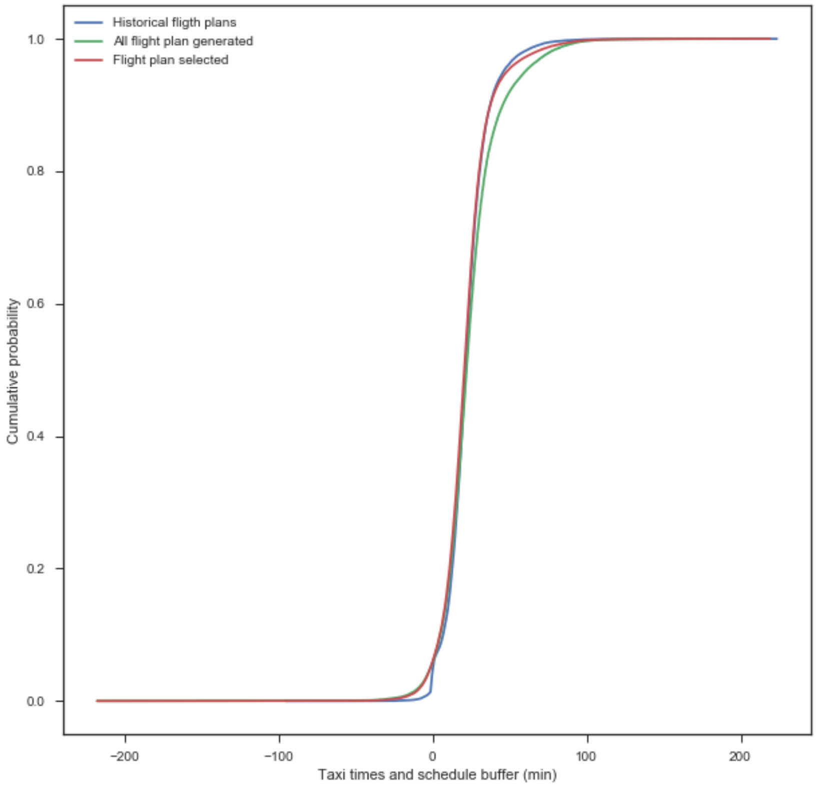

Figure 45. Cumulative probability of taxi and buffer time

Table 5. Quantiles and mean for FP taxi times and buffers (mins)

Type	Quantile						Mean
	0.1	0.2	0.5	0.7	0.8	0.9	
All flight plans generated	5	12	22	29	35	46	24
Flight plans selected	5	11	21	27	31	38	22
Historical flight plans	7	13	21	27	31	37	22

As shown in Figure 44 the times generated for the flight plan selected are very similar to the historical ones. This is further observed in Figure 45 and in Table 5, the distribution of taxi and schedule times are very similar for the flight plan selected and for the historical flight plans. Once again, if we compare with all the flight plans generated we can see that there is a higher discrepancy with respect to the historical flight plans. This is once again linked with the fact that longer flights have, in general longer buffers and there are more possible routes for longer flights and hence more observations of longer buffers, as shown in Figure 45.

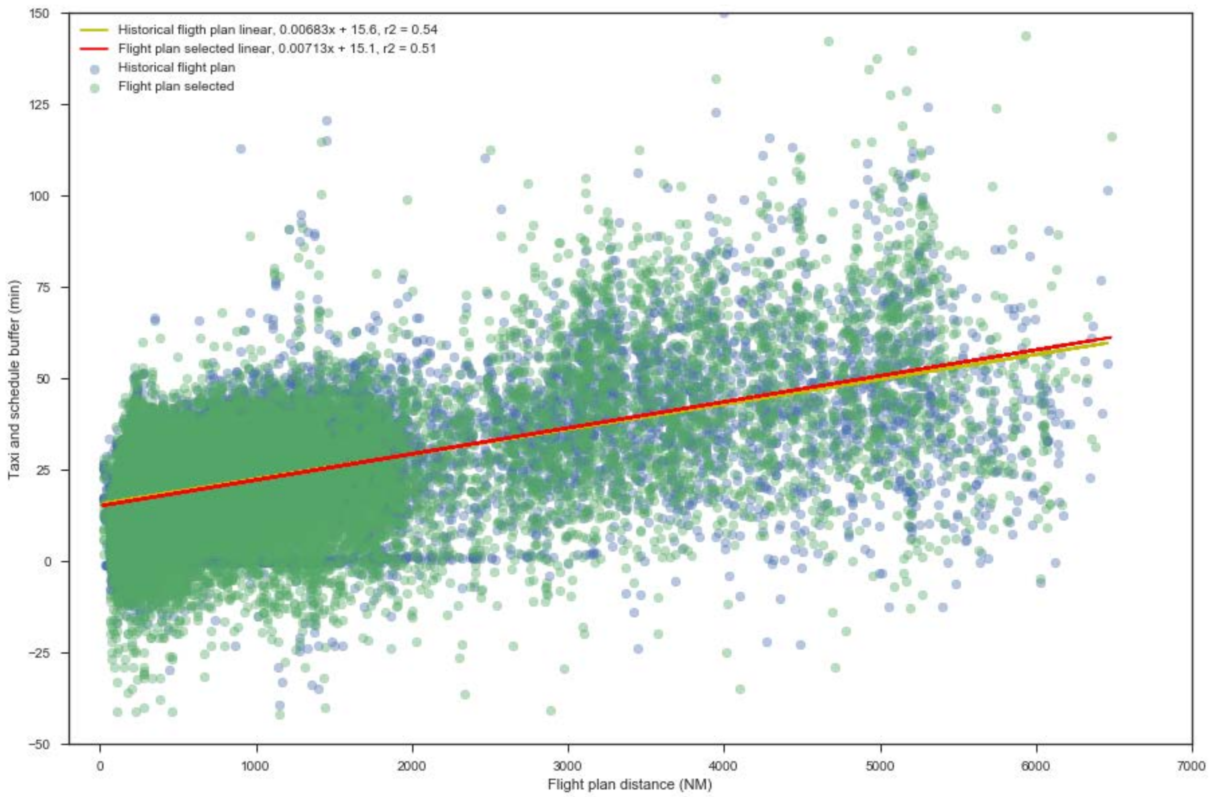

Figure 46. Taxi and buffer times as a function of FP distance (NM)

Note that in the historical data there are more with a buffer close zero that in the generated data, i.e. the duration of the flight plan is exactly the block time (5.6% of all flights in the historical data have a buffer in the -2 to 2 minutes while only 3.0% are in that situation).

(iii) Fuel usage

In the DDR dataset there is no information on the fuel usage by the aircraft so it is no possible to directly compare the historical fuel usage with the one obtained from the flight plan generator.

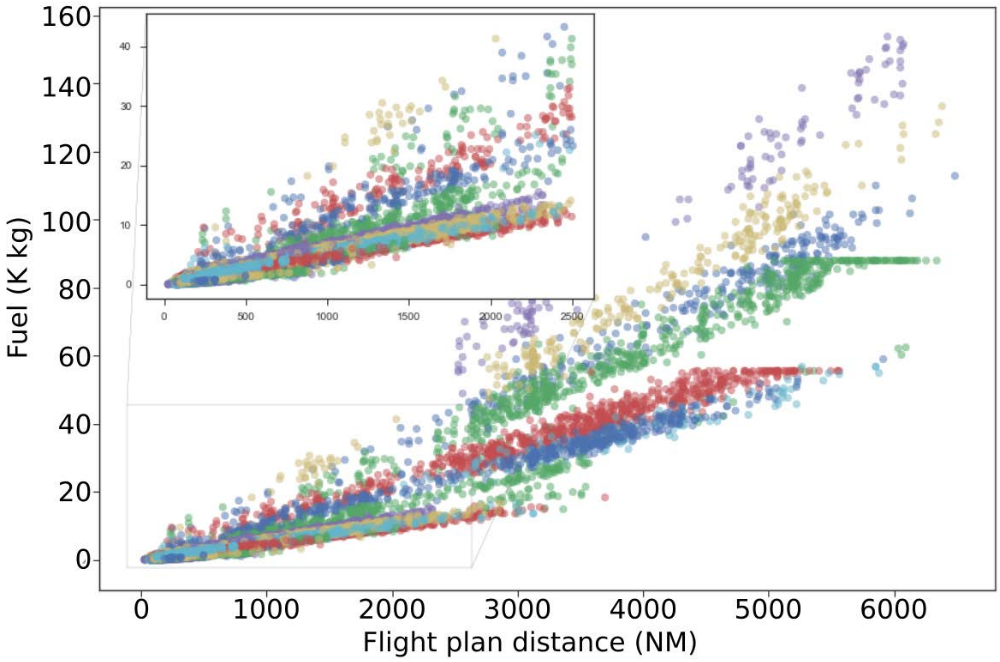

a) all aircraft types

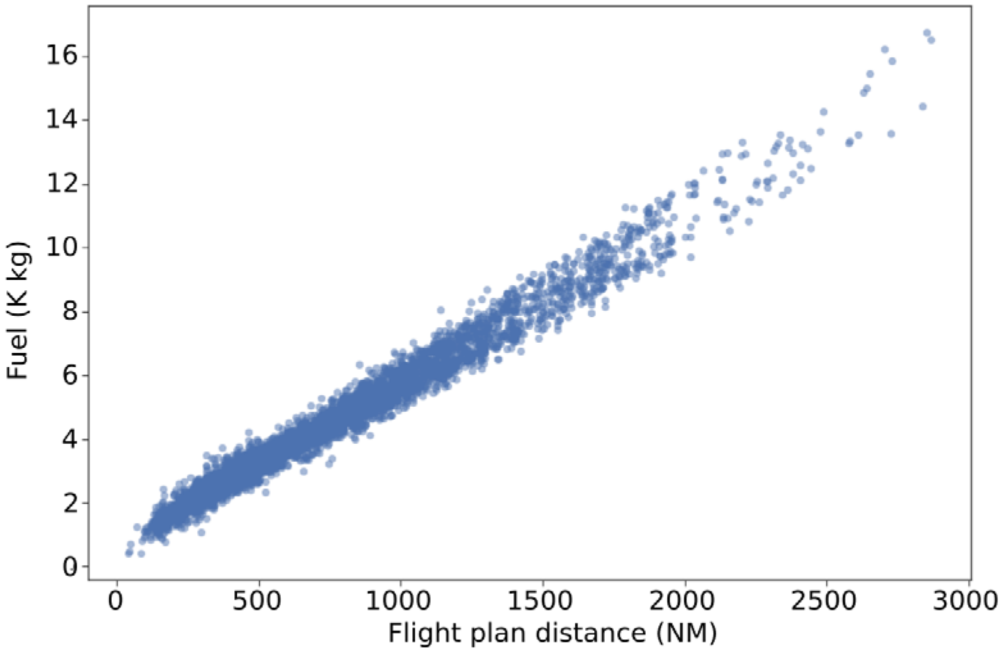

b) medium range twin engine aircraft type

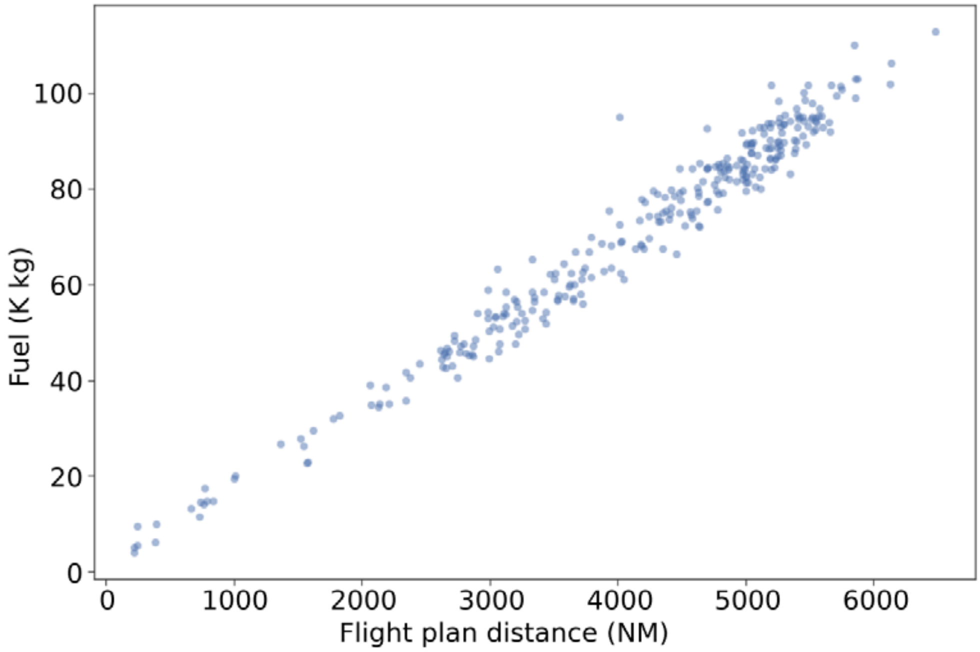

c) long range four engine aircraft type
Figure 47. Fuel as a function of FP distance (NM) per aircraft type

Figure 47 shows for different aircraft types the fuel used in the flight plan as a function of the flight plan distance. As expected there is a relationship between the flight plan length and the total fuel required. It is also interesting to observe how different aircraft type present different fuel increments as a function of flight plan distance, this is not only due to flight performance efficiency but also to the amount of payload they are able to transport.

(iv) En-route airspace costs

The en-route air service costs are dependent on the route selected, the airspace used and the MTOW of the aircraft. The relationship between distance and costs is not as clear as with fuel, the reason is that we are only computing the cost of the CRCO charges for the European airspace and adjacent NAS as explained in Section 2.2.1.3. This means, that very long flights might have low en-route air usage cost if they perform most of their flight outside areas we are considering (e.g., flights from Asia or from South America). This can be observed in Figure 48 where for long flights (greater than 3000 NM) there is a lower cost on charges than for shorter ones. See in Figure 48.b and Figure 48.c how the relationship between distance and air service cost is clearer for a smaller aircraft performing shorter routes. There is however, a high variability due to the different cost of airspace modelled.

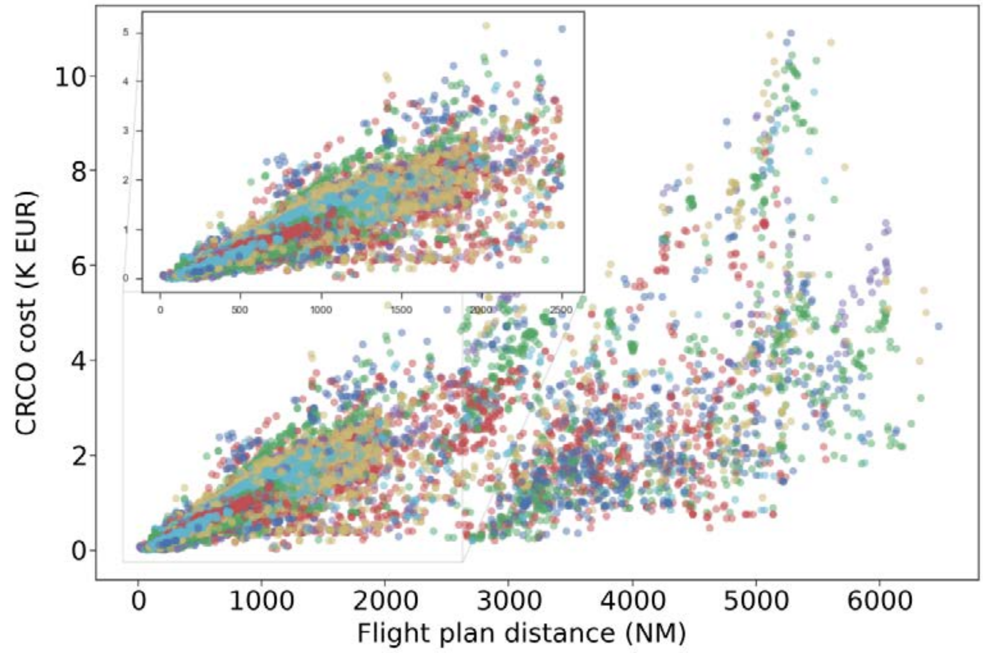

a) all aircraft types

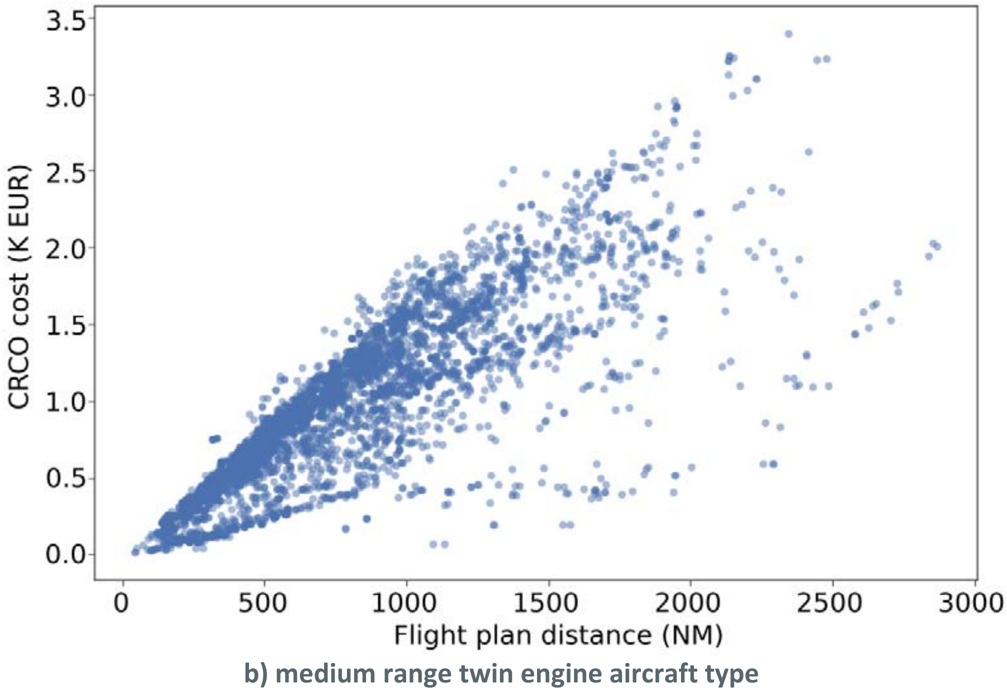

b) medium range twin engine aircraft type

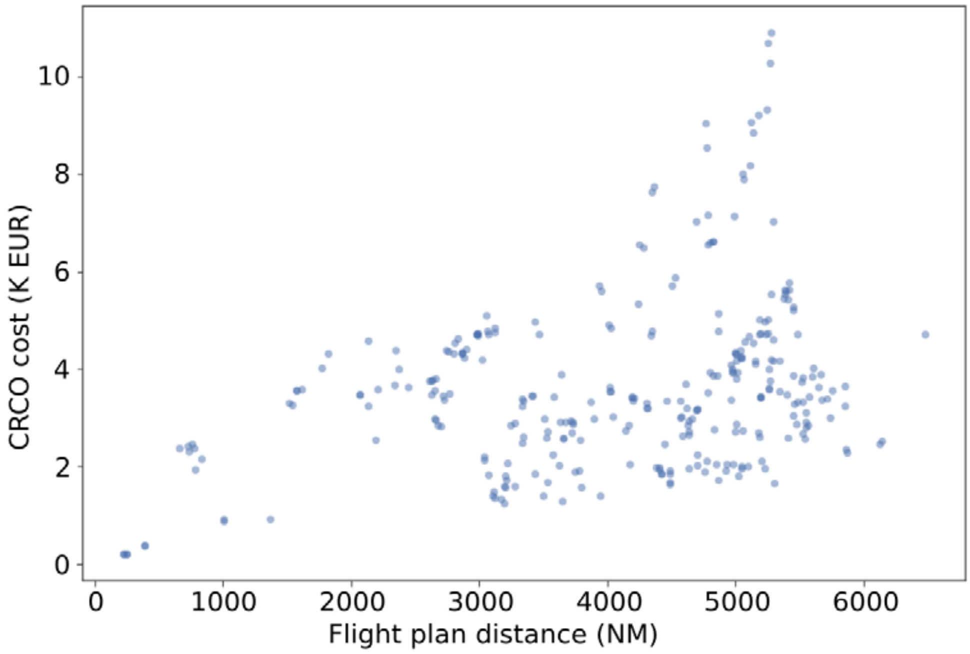

c) long range four engine aircraft type

Figure 48. En-route air service cost as a function of FP distance (NM) per aircraft type

As with the fuel, there is no information on the en-route airspace charges for the historical data. Figure 49 shows the revenue on en-route airspace charges per NAS.

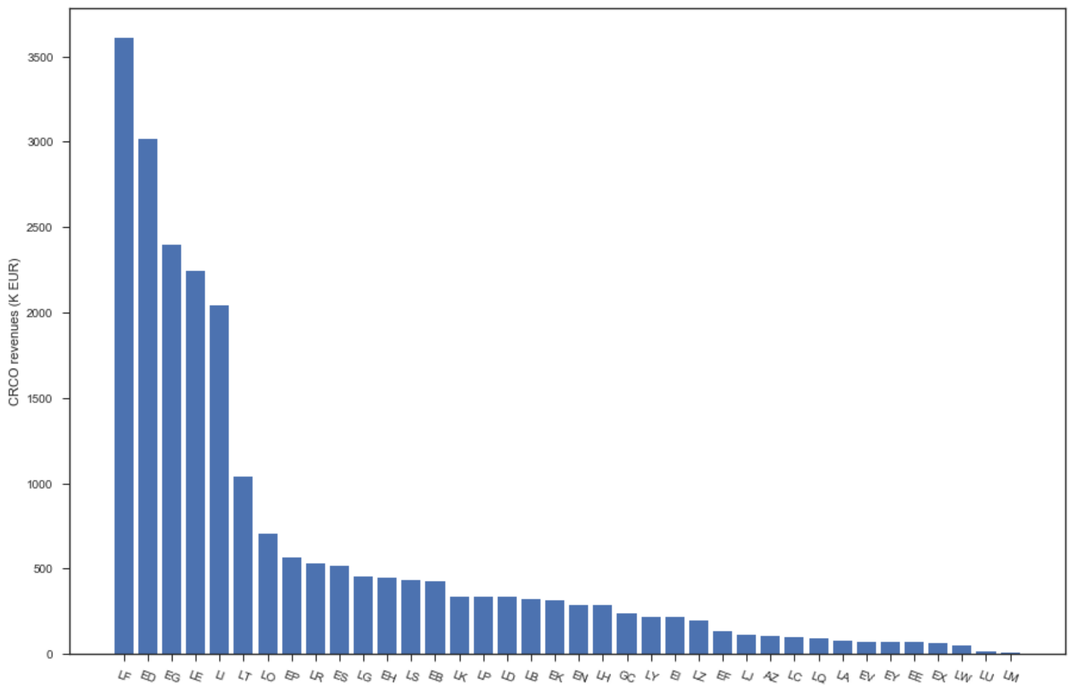

Figure 49. En-route air services revenue per NAS

Note EX: Shanwick airspace, GC: Spain Canary Islands, AZ: Portugal Azores.

In order to assess if the values obtained are in the right order of magnitude a comparison is done with reported data. EUROCONTROL's CRCO office publishes monthly en-route service units per NAS via the En Route Service Units Monitoring [11]. With this information it is possible to obtain the actual service units that were provided per NAS during 2016. The historical schedules we are using to generate the baseline traffic are from 12SEP14, for this reason, the service units of SEP16 are used. Using the same unit rates as those in the model (FEB18) we can compute the revenues for the SEP16 month and dividing it by 30 days we obtain an average revenue per NAS per day that is compared to the value obtained from the flight plan generator as shown in Figure 50. Figure 51 shows the difference in service units which consider aircraft distance and size within the different NAS.

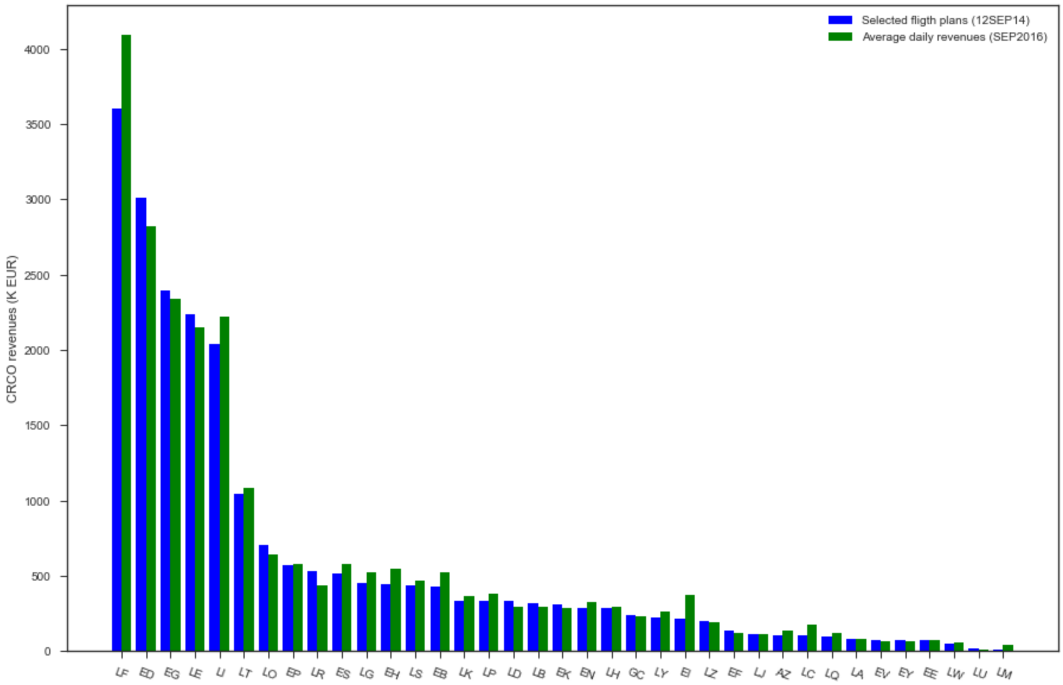

Figure 50. En-route revenues per NAS from FPs generated and reported service units

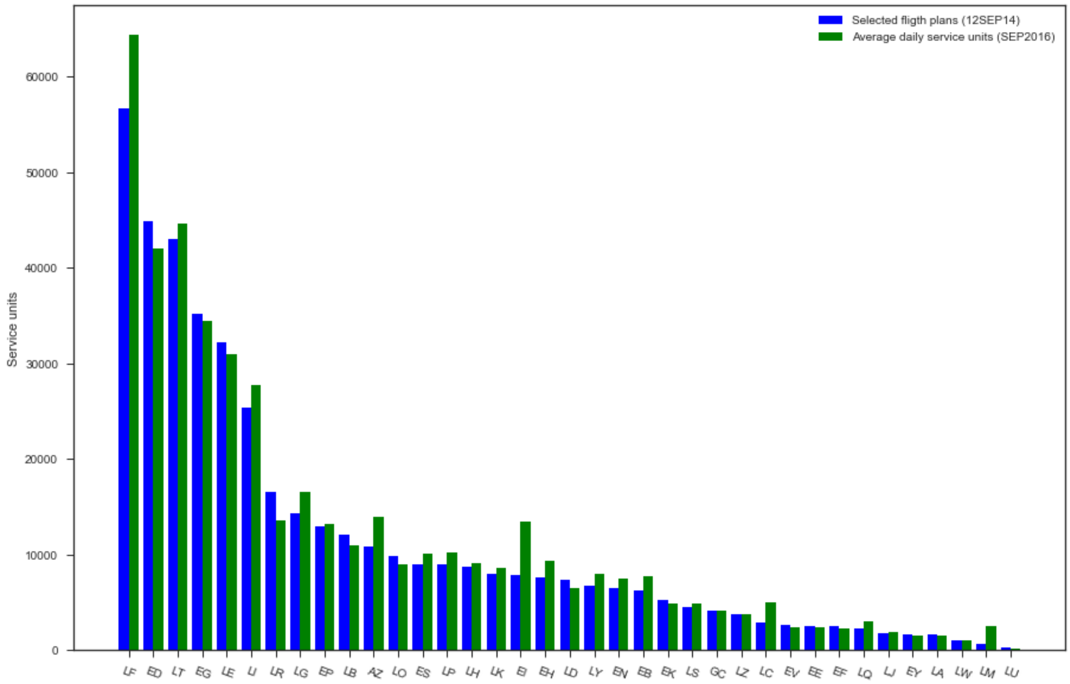

Figure 51. En-route service units per NAS from FPs generated and reported service units

The average error between the revenues obtained from the model for the day and the revenues computed based on the service units reported in 2016 is of 4.2% and the average difference in service units is 5.7%. Note that in general the estimation from the service units reported is higher, this is expected as the 2016 traffic is higher than 2014 and in the historical traffic used only passenger commercial traffic is considered. The revenues are linked to the service units which are related to the demand on the different airspaces, this allows us to see how the reported demand and the generated through the flight plans is similar.

3.2.2 Air Traffic Flow Management regulation generator

3.2.2.1 Calibration principles for the ATFM regulation generator

As described in Section 2.2.2, the ATFM generator is based on information on historical ATFM regulations and on the demand produced by the flight plan generator. An analysis of historical regulations and delay is needed to adjust the probability of being delayed and the intensity of this delay based on the historical data. The data analysed is primary sourced from DDR (AIRACS 1313-1413).

3.2.2.2 Input data for ATFM regulation generator

The input data for the ATFM regulation generator model calibration are the regulations and delays from historical data. All delays in the period AIRAC1313-AIRAC1413 have been analysed. The days when industrial actions were produced are filtered out to avoid unrealistic ATFM delay distributions. Only traffic departing from a country which is part of the centralised ATFM system in Europe, which can have delay assigned due to ATFM [18], is considered in this analysis, as other traffic will be excluded. The remaining traffic has been divided between delay due to capacity issues and 'other'.

This has provided us with the probability of having a flight crossing a regulation which is the most penalising for the flight (8.98%); this probability can be divided between the probability of being delayed due to non-capacity related regulations (5.09%) and the probability of being delayed due to capacity related regulations (3.89%). Three cumulative probabilities distributions of the amount of delay assigned can be computed: one for the total ATFM delay, one for the delay due to non-capacity, and one for the delay due to capacity.

The analysis of a week of traffic has allowed us to compute the demand and the probability of having delay for each ANSP due to capacity based on the regulations set that week. For each ANSP, the daily traffic increases at the beginning of the day and decreases at the night. The probability of having a regulation due to capacity should only be used when the traffic starts to be significant at the ANSP and should finish when it drops at night. For each ANSP, the demand of traffic considering the historical data of a week has been computed, then the time when the traffic volume (aggregated in 30-minute slots) is higher than the 40th percentile of the maximum observed in the day is computed defining in this manner the start time and end time when regulations due to capacity are possible per ANSP. Table 6 presents these start and end times when regulations are possible for some ANSPs and Figure 52 shows the demand variation over time for the days of analysis for a given ANSP. The limited number of days analysed might have an impact on the calibration of the ATFM delay probabilities and larger time series could be considered in future.

Table 6. Starting and ending time for ATFM regulations due to capacity for some ANSPs

ANSP	Starting time	End time
ED	05:15	20:15
EG	05:45	19:45
EP	04:45	20:45
LE	06:15	20:45
LF	06:15	20:15
LI	05:15	20:15

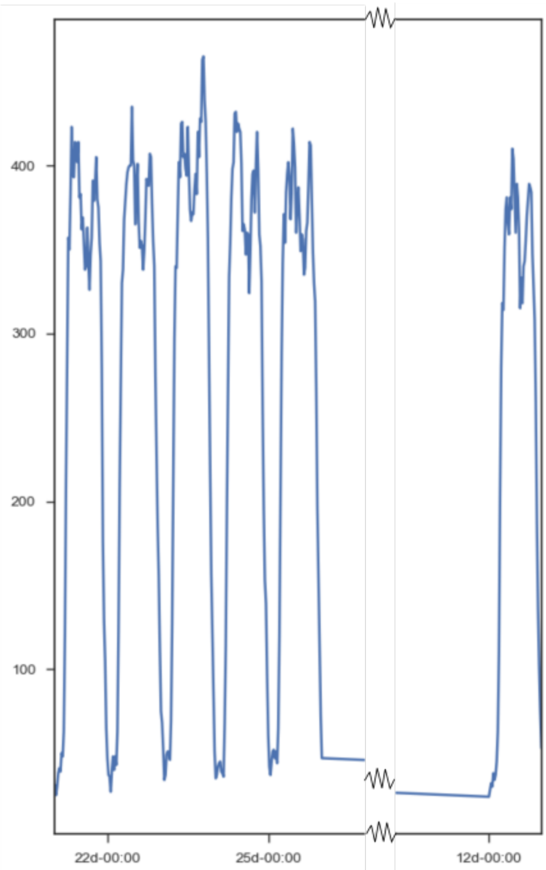

Figure 52. Historical traffic demand in LE for 6 days of analysis

In summary:

- Delay due to non-capacity regulations is based on historical analysis with a probability of having delay due to this and a cumulative probability function to generate the amount of delay. This delay generation is data driven and the delay generated will follow the same probability and distribution as the historical data.
- Delay due to capacity regulations, in this case, is from the analysis of a year of data. This cumulative distribution is used to generate the delay, but the probability of having delay assigned is disaggregated per ANSP. This means that the total probability after the assignment of delay is constructed from the traffic and the individual probabilities of assigning delay for each ANSP. Moreover, delay in the ANSPs due to capacity will be assigned only for flights crossing the ANSPs between the start and end time computed form the analysis of historical data.

In post-calibration, the probabilities of assigning delay need to be adjusted so that the total delay generated is close to the actual observed value in the current scenario.

3.2.2.3 Post-calibration of ATFM generator

The objective of the post-calibration is to ensure that the amount of delay generated is as close as possible to the current values in the current scenario. The same module to generate the delay as the one used in the tactical layer has been used to generate delay for the scenario of 2014 traffic. To the traffic of this scenario, the ATFM delay has been assigned 50 times computing the statistics of the results obtained and comparing them with the target one which comes from the analysis of the year of data. For more information on how delay is generated see Section 2.3.1.5.

As the probabilities of having delay per ANSP are based only on the analysis of a week of traffic, they have been adjusted with different calibrations. Four calibration options have been tested. The fourth one is selected for computations of the scenarios in Vista.

Table 7. Calibrations tested for ATFM probability of regulation due to capacity per ANSP

Calibration	Rationale	Probability non-capacity (5.09%)	Probability capacity (3.89%)	Probability all (8.98%)
0	Probabilities as estimated from the analysis of the week of traffic	5.11%	3.36%	8.48%
1	Countries with probability of zero modified to 0.1%	5.07%	3.45%	8.56%
2	Countries with probability of zero modified to minimum probability seen in week of data (0.153%)	5.07%	3.54%	8.61%
3	Countries with zero modified to 0.1% except countries which have had regulations due to capacity in the year of data which get the minimum seen (0.153%)	5.08%	3.52%	8.59%
4	First countries which have seen regulations due to capacity in the year but do not have probability now get minimum observed (0.153%). Then all countries which have a probability lower than 1% increment their probability by 15%. Finally, countries without probability get minimum observed (0.153%)	5.10%	3.62%	8.72%

Table 7 presents the four calibrations tested and the probabilities obtained when applying them on the 2014 scenario 50 times. The first calibration is to use the probabilities obtained directly from the analysis of the week of traffic. This has the disadvantage that some ANSPs did not report any regulation due to capacity on the days analysed. Therefore, it is not surprising that the probability of delay due to capacity across all the flights in Europe is of 3.36% in this case instead of the target 3.89%.

The next step, calibration 1, was to assign a non-null probability (0.1%) to the countries which did not have any delay. This, however, still produced lower amounts of delay than the target. In the following calibration, calibration 2, we tested assigning to the countries without probability the minimum probability observed for the countries that had probability (0.153%), this value is slightly higher than the 0.1% used in calibration 1 and therefore the results are better (3.54%), but still low.

Finally, we acknowledged that from the countries that have a probability of zero (because they didn't issue a regulation during the period analysed), some of them (EN, LJ, LY, ES, LH, EB, GM, LO and LK) have issued at least a regulation due to capacity issues in the year of analysis (AIRAC1313-1414) while the others have not. Therefore, we tested to assign to those countries the minimum probability observed and a non-null probability for the rest. Logically, this calibration 3, is assigning 0.1% probabilities to some countries which had a probability of 0.153% on calibration 2 and as a consequence the amount of delay generated is lower (3.52%), even if not significantly so, as the number of countries without probability is low, as is their traffic.

With the previous principle of having a larger probability for countries which have issued regulations in the past with respect to countries that have not, and considering that the overall amount of delay needed to be higher, although a few countries already had a probability which was particularly high (for example EP has a probability of 6.52%), we designed the final calibration. Calibration 4 assigns the minimum probability observed (0.153%) to the ANSPs that had a probability of zero but that issued regulations during the year. This increases by 15% (relatively) the probabilities for all the ANSPs which had a probability lower than 1% and finally assigned a non-null probability (0.1%) to the remaining countries without regulation. This ensures that the countries that have issued regulations have a slightly larger probability of assigning delay due to capacity; countries with low probability increase their values while remaining the high probability ones, and all countries have a non-null probability. This last point is important as these calibrated probabilities are the base (P_0) for the variation of the probabilities due to the demand and capacity changes in the different scenarios.

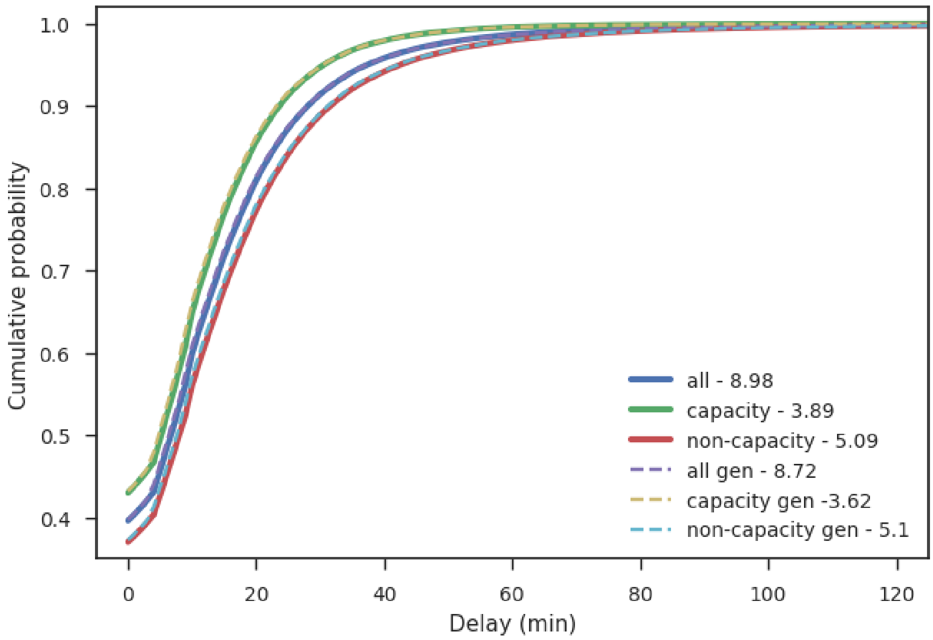

Figure 53. Probability distribution of ATFM delay with calibration 4

Figure 53 presents the cumulative distribution function for the ATFM delay for the analysis of the year of traffic and for the delay generated with calibration 4. Note that even if the flight goes through a regulation there is still around 40% probability of having a delay of zero minutes assigned.

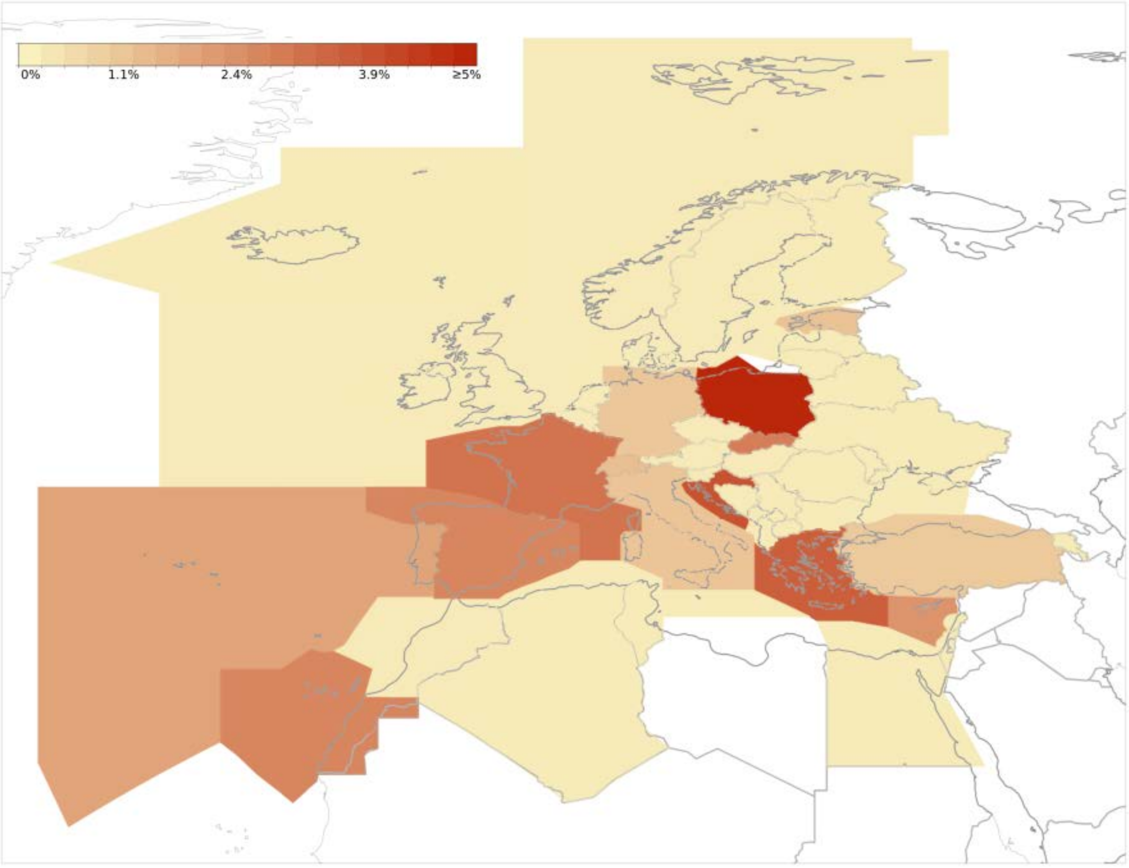

Figure 54. Probability having delay assigned due to ATFM capacity per ANSP in 2014 scenario

Figure 53 presents the probabilities of assigning ATFM due to capacity reasons per ANSP for the 2014 scenario with the final calibration selected.

3.2.3 Passengers itineraries generator

3.2.3.1 Calibration principles for the itineraries generator

The objective of the passengers itineraries generator is to transform flows into individual passenger itineraries. The assigned itineraries use information from different data sources to ensure that the minimum connecting times are respected at the airports. It is also important to validate that the percentage of connecting passenger at the different hubs is as expected from historical data and that the load factors on the flights is also aligned with industry reported values.

3.2.3.2 Input data for the itineraries generator

As explained in Section 2.2.3 the load factors of the flight have been adjusted to meet the values that are commonly reported by airlines. The results obtained from the assignment of passengers for the 2014 baseline scenario have been compared to the itineraries used in previous projects (e.g. SESAR WP-E, 'ComplexityCosts' [19]) which represent a typical day of operations in Europe. Those historically generated itineraries were validated (*ibid*.) against GDS, ACI Europe and airline data.

3.2.3.3 Post-calibration of itineraries generator

There are two executions of the strategic layer producing two sets of schedules and passenger flows per scenario. Therefore, the itinerary generator has been executed twice with the baseline 2014 traffic and passengers flows.

The historical traffic data of ComplexityCosts had a total number of passengers of 3.4M passengers, both executions of the Vista pre-tactical layer have also produced 3.4M passengers, as shown in Table 8. Note also that the number of passengers with connections is also similar even if slightly lower than in the ComplexityCosts data. This correspondence in the data also validates the outcome of the strategic layer with the passenger flows.

Table 8. Passenger itineraries comparison with 2014 itineraries

Source	Passenger type	Number of passengers ('000)	Passengers with 1 connection ('000)	Passengers with 2 connections ('000)	Number of different itineraries
ComplexityCosts 2014 itineraries	Standard	3 151 (92.3%)	284 (8.3%)	10 (0.3%)	7158
	Premium	264 (7.7%)			
	Total	3 415 (100%)			
Execution 1 2014 baseline	Standard	3 197 (94.0%)	276 (8.1%)	9 (0.3%)	3241
	Premium	205 (6.0%)			
	Total	3 402 (100%)			
Execution 2 2014 baseline	Standard	3 193 (93.9%)	275 (8.1%)	9 (0.3%)	3200
	Premium	207 (6.1%)			
	Total	3 400 (100%)			

As explained in Section 2.2.3, the passenger itinerary generator also defines how many passengers are 'premium' and 'standard'. This is done with a probability depending on the difference in price with respect to the fitting of prices as a function of total trip distance. This value has been adjusted so that the number of premium passengers generated is consistent with that observed in the historical data, as shown in Table 8.

The indicator that shows the greater difference is the number of different itineraries that are observed per scenario. In the ComplexityCosts 2014 data there were a total of 7 158 different itineraries, i.e., different flight routes taken by passengers, while in the Vista generated itineraries the number is reduced to only 3 200. This could be due to the fact that the optimisation used to generate the itineraries from the flows is maximising the number of passengers that are set in the flight but is not distributing the passengers. For example, if a given flow has ten options (i.e., ten possible itineraries to be taken) the algorithm is simply trying to maximise the number of passengers from the flow that can travel, not the number of options used, and it could be that all the maximum numbers of passengers assigned can be achieved using only two of the options, instead of distributing the number of passengers across all the possible options. This is something that could be improved in the future by adding these considerations in the objective function of the passenger itineraries generator. However, even if fewer combinations are used, the number of passengers connecting and the load factors of the flights are consistent with historical data.

3.3 Tactical layer

The calibration of the tactical model is done using some of the data sources identified in D2.1. Those data sources relate to 12SEP14. Most of the calibrations were already performed in the previous iteration. However, the model has to be calibrated again to include the new sub-models added in this final iteration. This document includes the whole calibration process to be self-explanatory.

3.3.1 Gate-to-gate simulation

The gate-to-gate simulation model includes several submodules that had to be calibrated in order to produce valid results. Most of them were already calibrated, either in the previous iteration or in previous projects. This section identifies the most important ones and describes the calibration processes in different levels of detail, because some of them are just input variables from upper layers.

3.3.1.1 Flight cancellations

Operational cancellations in 2014 remained stable at 1.5%, with peaks of around 8% on days with industrial action – as shown in Figure 55, these days correspond to airline strikes [24].

Figure 56 presents the average percentage of cancellations for each month in 2014. The peaks observed in April (1.9%) and October (1.8%) are due to strikes at different airlines. In December, there were multiple ATC industrial actions combined with a technical problem at Heathrow, which increased the cancellation rate to 2.4% (*ibid*.).

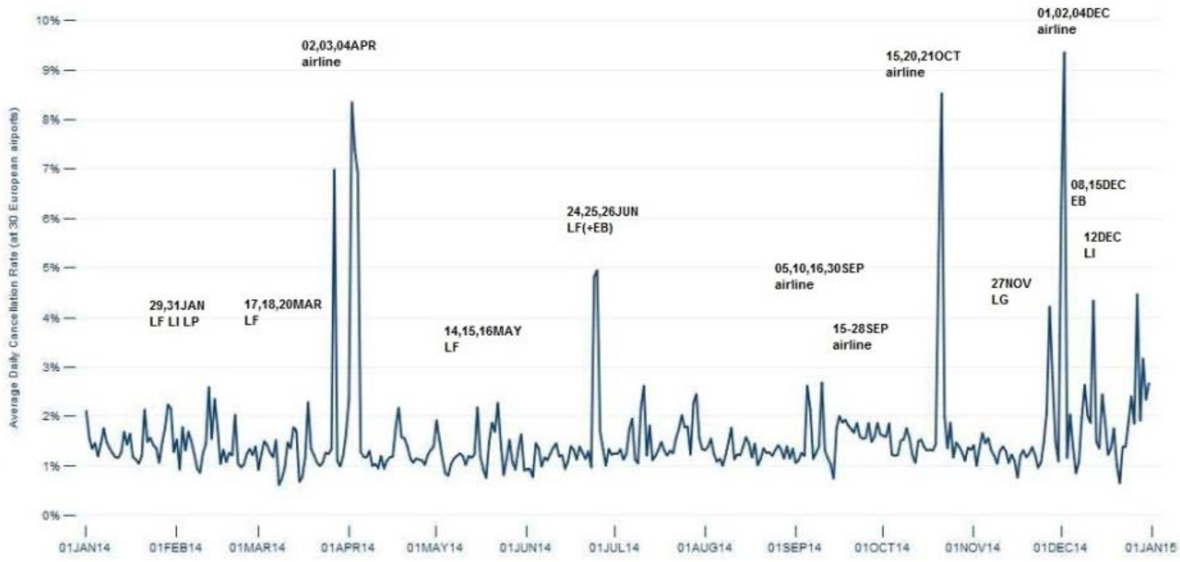

Figure 55. Average daily cancellations 2014 [24]

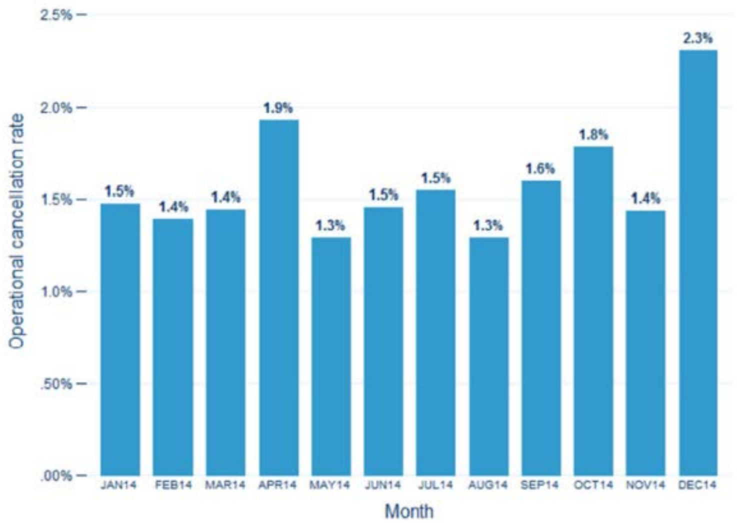

Figure 56. Monthly share of operational cancellations 2014 [24]

Figure 55 shows the cancellation rate experienced on a daily basis, industrial actions are marked at some peaks. In May 2014, French ATC industrial action generated peaks at 2% of cancellations and the industrial actions of 24-26 June increased the cancellation rate up to 5% [24]. Traffic data on days where industrial actions were implemented (30JAN14, 15MAY14, 24JUN14, 19MAY16, 26MAY16) have been analysed in detail. A comparison between the total number of flights that operated in the ECAC area with respect to the preceding and following weeks, shows that there is a reduction in the total ECAC traffic ranging between 0.4% and 3.9%. Mindful of all of these considerations, a rate of between 2% and 5% of cancellations is considered for flights affected by ANSP industrial action (note that there is a minimum, baseline level of 1.5% cancellations on a nominal day, for all the flights operating in the ECAC area).

3.3.1.2 Minimum turnaround time

The objective of this data analysis is to determine the minimum turnaround times (MTTs) to feed the Vista tactical layer. The MTT was defined to be approximated as the 2nd percentile of one month of historical data of turnaround times, provided the sample is large enough to estimate the 2nd percentile value with 95% confidence. Table 9 shows the final values used in the model, as well as an approximation for the turnaround times distribution in each category, whenever the data sample is large enough. These values, distributions, errors estimation and confidences were obtained following the methodology detailed below. (Additional turnaround distributions that have not been used in Vista are also shown.)

Table 9. Minimum turnaround times used in Vista

<table>
<tr><th>airline</th><th>aircraft type and turnaround times</th><th>airport</th><th>MTT</th></tr>
<tr><td rowspan="3">all</td><td rowspan="3">Heavy aircraft
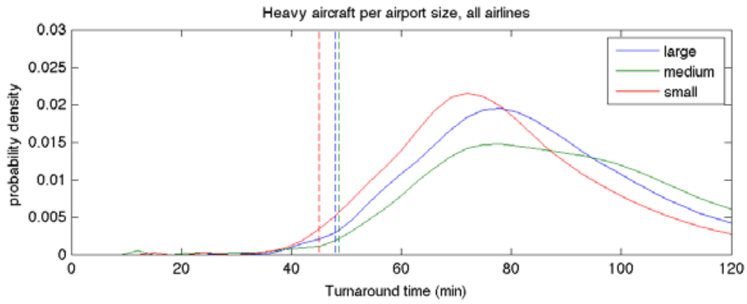
</td><td>large</td><td>45’11”</td></tr>
<tr><td>medium</td><td>48’36”</td></tr>
<tr><td>small</td><td>47’59”</td></tr>
<tr><td rowspan="3">REG</td><td rowspan="3">Medium aircraft
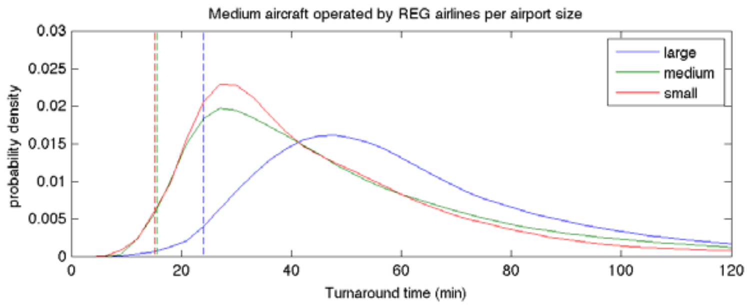
</td><td>large</td><td>24’02”</td></tr>
<tr><td>medium</td><td>15’35”</td></tr>
<tr><td>small</td><td>15’13”</td></tr>
</table>

<table>
<tr><th>airline</th><th>aircraft type and turnaround times</th><th>airport</th><th>MTT</th></tr>
<tr><td rowspan="3">CHT</td><td rowspan="3">Medium aircraft
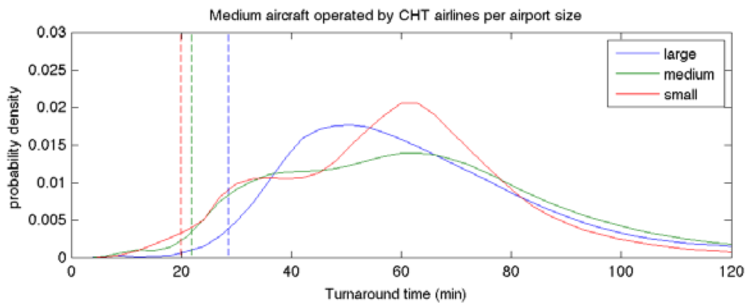
</td><td>large</td><td>28’38”</td></tr>
<tr><td>medium</td><td>21’57”</td></tr>
<tr><td>small</td><td>19’58”</td></tr>
<tr><td rowspan="3">LCC</td><td rowspan="3">Medium aircraft
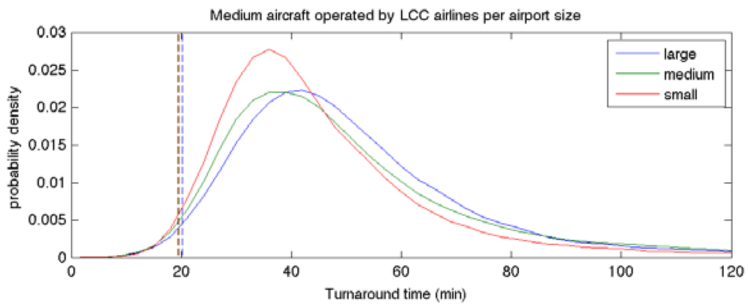
</td><td>large</td><td>20’47”</td></tr>
<tr><td>medium</td><td>20’11”</td></tr>
<tr><td>small</td><td>19’42”</td></tr>
<tr><td rowspan="3">FSC</td><td rowspan="3">Medium aircraft
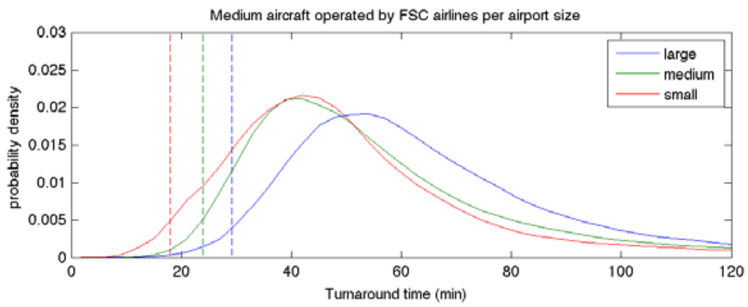
</td><td>large</td><td>29’09”</td></tr>
<tr><td>medium</td><td>23’57”</td></tr>
<tr><td>small</td><td>17’57”</td></tr>
<tr><td>all</td><td>Light aircraft
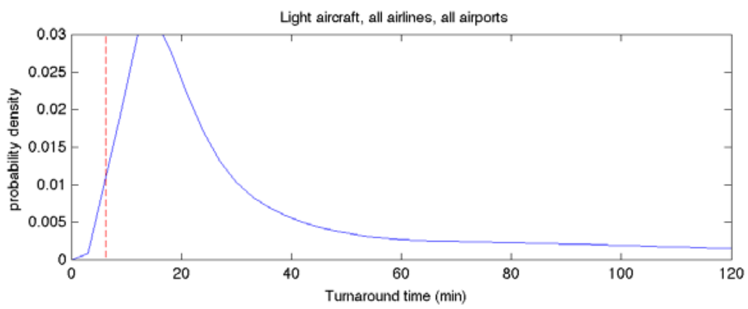
</td><td>all</td><td>6’17”</td></tr>
</table>

Table 10. Additional turnaround times distributions (not used in Vista)

description	turnaround times	filter	MTT
per airport size all aircraft all airlines	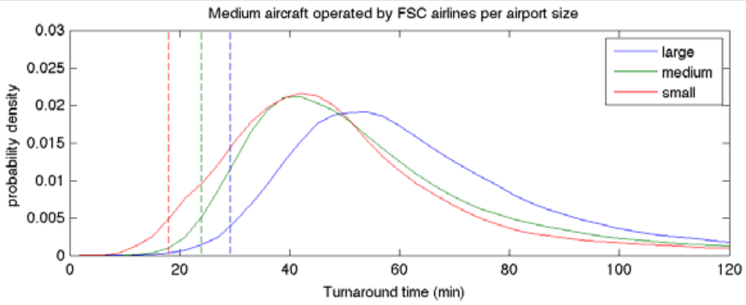	large airport	25’24”
		medium airport	21’33”
		small airport	18’03”
per aircraft weight all airports all airlines	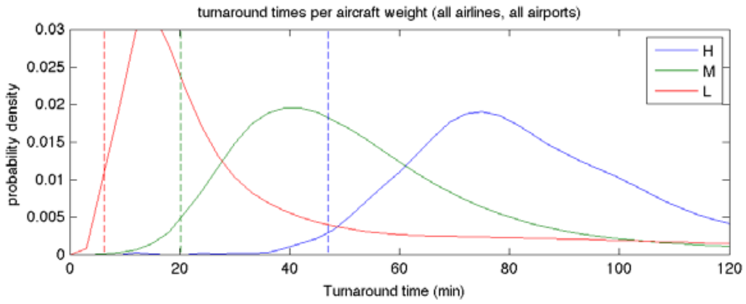	heavy aircraft	46’55”
		medium aircraft	20’07”
		light aircraft	6’17”
per airline type all airports all aircraft	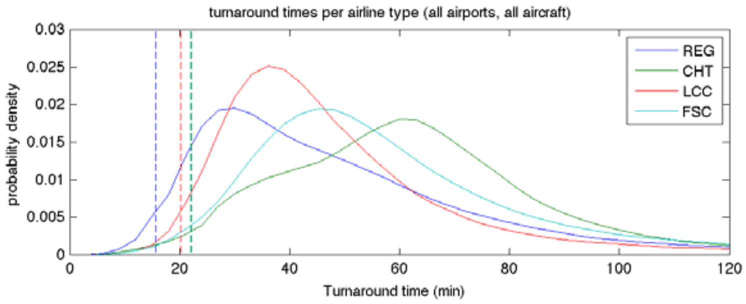	REG	15’36”
		CHT	22’03”
		LCC	20’07”
		FS	21’59”
overall all airlines all aircraft all airports	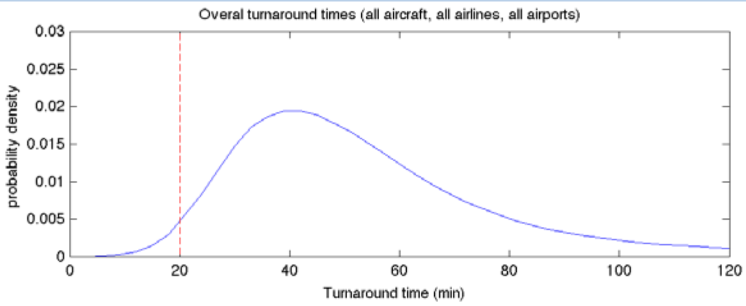		20’04”

Statistical method

The MTT is defined as the sample’s 2nd percentile, provided the sample is large enough to provide an error <2%, estimating also the confidence of the error. The statistical methods used to estimate the 2% value are as per:

1. Identify rotations in the traffic sample, using the aircraft registration, arrival and departure airports, and timing (OBT, duration and taxi-in).
2. For each rotation compute the turnaround time and classify the values by **airport size, aircraft size and airline type**.
3. Estimate the **minimum sample size** to ensure an error smaller than 2%:
   a. The probability distribution function of the sample is estimated using a Normal kernel smoother (no data censoring), with the 2% value selected.
   b. The p-value of likelihood of being Normally distributed with empirical mean and deviation using a simple Kolmogorov-Smirnov test.
   c. Assuming data are Normally distributed, use the confidence interval for proportions to determine the **minimum sample size** necessary to produce an error smaller than 2%, with a confidence of over 95%.
4. If the sample size is not large enough, drop one or more categories in the following order: airline type, airport size and aircraft weight.
5. Repeat above steps, as required.

The uncertainty in actual turnaround times is modelled as an Exponential distribution characterised by $\lambda = 5$ minutes, truncated to 30 minutes (estimated from ALL_FT+ dataset).

3.3.1.3 Primary delay unrelated from ATFM

This layer distinguishes two reasons for primary delay: ATFM and non-ATFM delay. The first is heavily related to the flight plan selection, so it is covered below. The latter is model stochastically as follows:

- exponential distribution model characterised by $\lambda = 15$;
- distribution truncated to 90 minutes;
- chance of having ATFM delay = 15%;
- values were fitted with data from ALL_FT+ dataset.

3.3.1.4 Flight plan selection

Calibrated flight plan selections are offered as an input from the pre-tactical layer, as well as the list of options. Consequently, the baseline scenario just selects the best option offered by the pre-tactical layer.

3.3.1.5 Primary delay due to ATFM

The calibration of the delay due to ATFM is explained in detail in Section 3.2.2. How the mechanism works to assign the delay is explained in Section 2.3.1.5.

3.3.1.6 Minimum connecting time

The minimum connecting time (MCT) is the shortest time interval required for a passenger/baggage to connect between flights at an airport. Standard MCTs are available for connections between flights that serve: domestic to domestic; domestic to international; international to domestic; and international to international. A single MCT per airport is required by the model, and is used with all connection types. Standard MCTs have been purchased from Innovata LLC who maintain the schedule reference service (SRS) database on behalf of IATA.

An average standard MCT has been calculated for each of the 200 ECAC airports in scope, ranging from 14 minutes at Oulu Airport to 3 hours 45 minutes at Zvartnots Airport (Figure 57). Overall, the average MCT is 45 minutes.

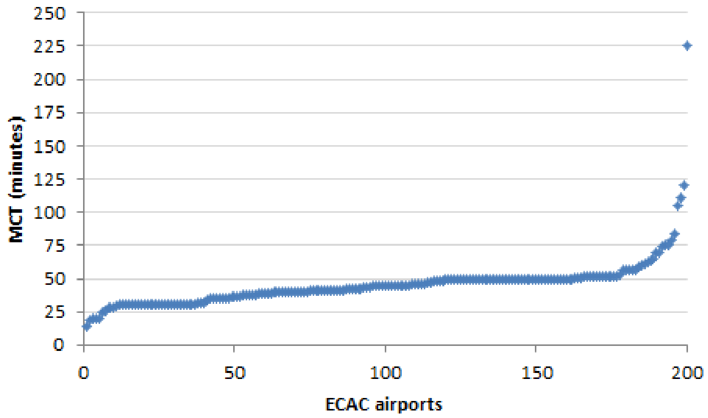

Figure 57. Minimum connecting time (MCT) for each ECAC airport

The sensitivity of the final model may be tested by increasing all the MCTs by 10%. During the execution, some uncertainty was added using a Normal distribution centred at 0, with $\sigma = 0.05$.

3.3.1.7 Additional delay for waiting connecting passengers

The model behaviour here depends on the passenger management tools, as described in Section 4.6.3.

3.3.1.8 Taxi times

The PRISME actual taxi-out times are used, although 'actual' times are often the same as 'filed' times for non-A-CDM airports (i.e. the vast majority). Additional taxi-out data have been sourced from CODA to supplement the PRISME data. For example, the mean taxi-out time at Madrid Barajas Airport in SEP10 was 19.70 minutes (standard deviation 8.24; variance 67.88) based on 16 242 departures.

For airports with missing 'actual' times, a stochastic element is incorporated based on summary statistics from the CODA data. This avoids the unrealistic assignment of filed taxi-out times to all aircraft.

For each airport and each carrier at that airport, a sample of 'nominal' weather and 'bad' weather days is taken. For each airport and carrier, a Normal distribution of taxi-out times is assigned for each type of day, and delays are drawn at random from that distribution, according to the type of day being modelled. Where such data are insufficient to produce a distribution (fewer than 100 observations for an airport-carrier combination), the airport nearest in size in terms of passenger numbers, which has sufficient observations, is used instead.

3.3.1.9 AMAN and DMAN capacity

AMAN and DMAN capacity is calibrated per airport and provided as an input from the strategic layer. To avoid assigning an extremely low capacity at some airports due to the fact that they were not operating at 100% capacity any time during the traffic sample period, a fixed capacity of 45 movements per hour will be assigned as the minimum possible capacity.

3.3.1.10 Passenger reaccommodation policy

Passengers are reaccommodated according to the base rule described here, and, alternatively, under the scenario rules described in the scenarios section. Scenario rules were applied to aircraft in the model to compare the comparative effectiveness of such rules in reducing the costs of delay to the airline, for example. The operational day is deemed to end at 0400 (local time) the following day, as a cut-off time for passenger reaccommodation.

The following sub-rules apply to the passenger reaccommodations:

- passengers are reaccommodated only onto flights where seats are available;
- passengers are reaccommodated according to the minimum number of changes involved in reaching their destination (e.g. they are booked onto direct flights, if possible, before indirect options are considered);
- no originally-booked passenger may be 'bumped' (replaced) by an attempted reaccommodation (although this may occur in practice);
- once a passenger has been reaccommodated, they may not be subsequently 'bumped' by a 'higher priority' passenger.

The following prioritisation order applies:

- full-service airline premium passengers;
- full-service airline standard passengers;
- all other passengers.

Table 11. Hierarchy of interlining

Carrier type	Rebooking onto next available flight according to departure delay of:		
	Up to 2 hours	2 – 5 hours	>5 hours
full-service airline premium passengers	any carrier	any carrier	any carrier
full-service airline premium passengers	booked/alliance only	any carrier	any carrier

Carrier type	Rebooking onto next available flight according to departure delay of:		
	Up to 2 hours	2 – 5 hours	>5 hours
full-service airline premium passengers	booked carrier only	booked carrier only	booked carrier only

Table 11 shows the rules for rebooking passengers according to the length of the expected delay at the final destination (regardless of length of haul, as a simplification), and the passenger type. "Booked/alliance only" means that the passenger is only rebooked onto the original carrier, or a co-member of an alliance network. Intra-carrier bookings take precedence over all other rules if any conflict arises.

If a passenger has not been reaccommodated by the above process, they are added to a wait list of unaccommodated passengers, one of which is held for each airport. For each flight departure, the reaccommodation rule is re-run for all passengers in the wait list after all missing connection passengers are rescheduled, as new reaccommodation opportunities may arise. The model assumes that only 20% of the passengers prefer to return to their point of origin. Passengers are reaccommodated onto aircraft regardless of ticket and cabin class, in that aircraft are simply occupied according to total seat space. No account is taken of the costs associated with upgrading or downgrading.

3.3.1.11 Climb uncertainty

In order to generate the uncertainties during the climb phase, the difference between the estimated time required from take-off to reaching FL180 according to the finally submitted flight plan and the actual time required to reach FL180 from departure for all the flights going to a hub during the period AIRAC 1313 to AIRAC 1413 (i.e., 12 December 2013 to 07 January 2015) have been analysed. By computing this climb time to FL180 the uncertainties due to climb and the departure TMA are considered. The differences in time are presented in Figure 58, which are approximated with the Normal distribution defined in Table 12.

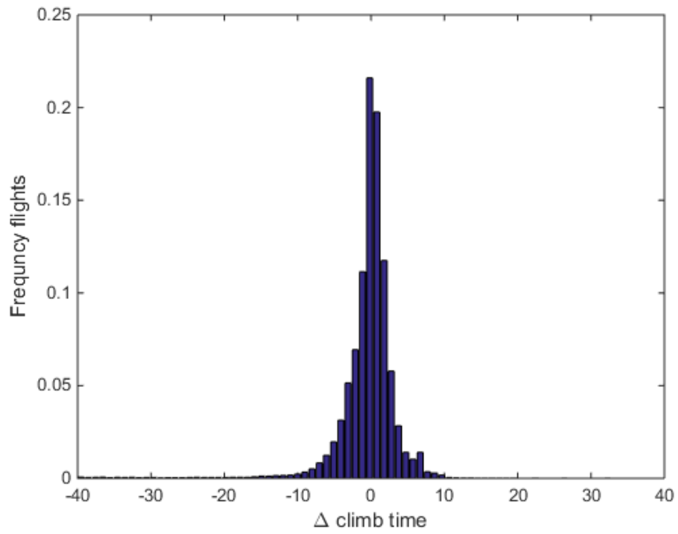

Figure 58. Time variability on climb phase at reaching FL180

Table 12. Variability on climb phase Normal distribution parameters

μ mean	σ standard deviation
-0.4	4.3

3.3.1.12 Cruise uncertainty

Following the same principle as in climb uncertainty, the estimated time from reaching FL180 in climbing until reaching FL180 on descend have been analysed. In this manner the distribution of time during the cruise has been estimated, as shown in Figure 59, which can be approximated with a Normal distribution of parameters (μ=-1.2 σ=6.5). This means that 95% of the values are in the range -14.3 to 11.7 min.

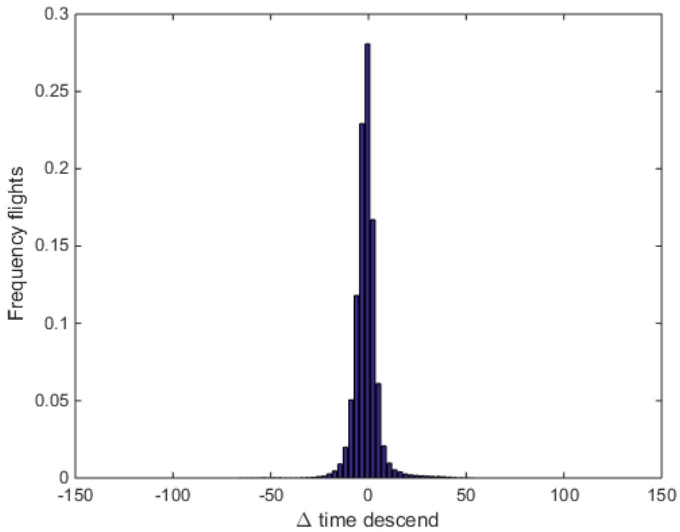

Figure 59. Difference between planned and actual time to cruise

Planned and actual time to cruise, from FL180 during climb until FL180 during descent.

This methodology has, however, some problems. This error in time accounts for uncertainties on weather and planned and actual route, but it also is affected by profile changes and tactical modifications of airline speed profiles. Therefore, to remove those uncertainties, the analysis of the distance between FL180 in climbing until FL180 in descend has been carried out, doing a comparison between the finally submitted flight plan and the actual flown for all the flights inbound to a hub in the AIRAC 1313 to AIRAC 1413 period. Figure 60 shows these differences in flight distance. Table 13 presents the parameters of the Normal distribution that is fitted to the data. This variability in distance will lead to an uncertainty in flight time for the cruise that will depend on the flight speed. As a comparison with the difference between planned and actual flight times, with this methodology, 95% of the values will be in the -49.6 to -25.2 NM, leading to a time variability of -6.6 to 3.4 min for a flight cruising at FL360 at M0.78.

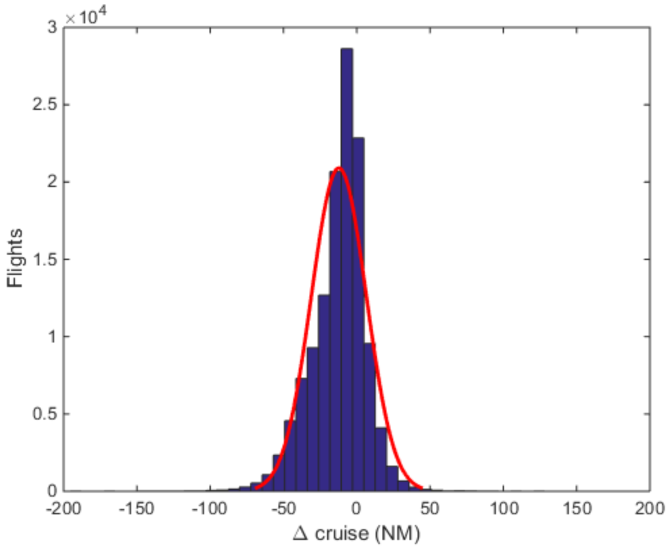

Figure 60. Differences in cruise distance between actual and planned

Differences in cruise distance between actual and planned from FL180 during climb until FL180 during descent.

Table 13. Variability on cruise phase Normal distribution parameters

μ mean	σ standard deviation	95% values (Time at FL360, M0.78)	
		Min	Max
-12.2	18.7	-49.6 NM (-6.6 min)	25.2 NM (3.4 min)

3.3.1.13 Delay recovery model

The delay recovery model is as described in Section 4.6.1.

3.3.1.14 Passenger delay costs (to the airline)

The passenger delay costs model is calibrated with data from previous and on-going research projects undertaken by UoW. Two different delay costs were considered:

- ‘hard’ costs, e.g., duty of care and compensation;
- ‘soft’ costs, e.g., estimation of loss of market share.

The duty of care costs are shown in Table 14 and Table 15, estimated by ‘low’, ‘base’ and ‘high’ cost scenarios.

Table 14. Costs applied duty of care as a function of aircraft operator and ticket

Aircraft operator (& ticket) type	Cost applied

Aircraft operator (& ticket) type	Cost applied
FSC (premium)	average of high and base
FSC (standard)	base
REG, LCC, CHT	average of low and base

Table 15. Costs (euros) due to duty of care, per passenger, by minutes of delay

Minutes of delay	High	Base	Low
0 – 90	0	0	0
90 – 120	2.2	1.8	0.0
120 – 180	10.2	8.4	5.1
180 – 300	25	21	13
300 – 480	28	23	14
480 – Overnight	109	90	55

Passenger regulations have evolved since the introduction of Regulation 261 in 2005 [20]. According to Regulation 261 and subsequent amendments [21], passengers are entitled to compensation when their flight is delayed on arrival. The following table defines the amount of money each passenger is entitled to.

Table 16. Cost due to compensation per passenger by minutes of delay

Flight type as a function of GCD	Delay in minutes		
	120 ≤ d < 180	180 ≤ d < 240	> 240
Short haul – distance < 1500 km	0	€250	€250
Medium haul – 1500 km < distance ≤ 3500 km	0	€400	€400
Long haul – distance > 3500 km	0	€300	€600

Based on consultation, the passenger uptake rate is 11, i.e., only 11% of passengers entitled claim the compensation. Flight distance is determined as the scheduled distance (GCD) from the first leg that is delayed at departure (at least 180 minutes) onwards. Table 16 presents the compensation passengers are entitled to as a function of the delay and flight type.

'Soft' costs manifest themselves in several ways. Due to a delay on one occasion, a passenger may (later) defect from an unpunctual airline as a result of dissatisfaction (although quite possibly reversing this defection subsequently). Such a loss may be considered to be largely the gain of another airline, gaining a passenger who has transferred their custom. When such scalable costs (multiplied over a period of time or a network) are assessed, only some net loss to the airlines is likely (for example due to trip mode substitution, trip consolidation, trip replacement (such as a teleconference) or trip cancellation). The passenger soft cost of delay is thus appropriately bounded.

For assigning the soft costs of delay, a logit function is used to describe passenger dissatisfaction (δ; normalised) against various levels of delay. This curve is used to distribute the soft cost as a function of delay duration, and may be thought of as a proxy for the propensity of a passenger to switch from a given airline, to some other choice, after trips with given delay experiences.

$$\text{Logit function: } \delta = \frac{1}{k\left(1+e^{a-bf}\right)} - k'$$

This is plotted in Figure 61 (black curve) and has the desirable characteristics of maintaining a low value for some time, then rapidly increasing through a zone of 'intolerance', before levelling off. Quantification of the saturation of delay inconvenience and crossovers in Kano customer satisfaction 'requirements' contributed towards this model [25]. Relationships between market share, punctuality and customer satisfaction were also examined.

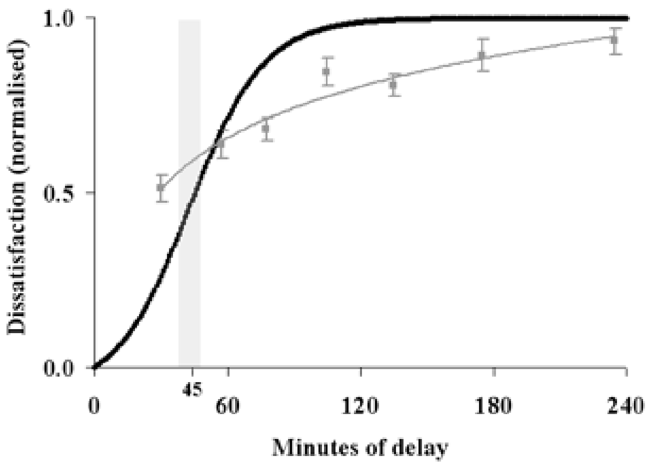

Figure 61. Passenger dissatisfaction as a function of delay duration [22]

Euro costs are assigned using δ as a weight; per-minute values are saturated by two hours:

Table 17. Passenger soft costs of delay per minute, by three cost scenarios

Scenario	Delay (min)								
	5	15	30	60	90	120	180	240	300
Low	0.01	0.02	0.07	0.19	0.25	0.27	0.27	0.27	0.27
Base	0.02	0.09	0.25	0.69	0.91	0.96	0.97	0.97	0.97
High	0.03	0.10	0.28	0.77	1.01	1.06	1.08	1.08	1.08

Note that the values in the previous table are per minute and per passenger. The soft cost estimate at any given time for a flight, is based on the sum of the individual passenger's delays at their final destination. Costs per minute are extrapolated/interpolated (to the nearest minute), from the values in the table, according to the cost scenario being used. The soft cost assignment is deterministic and based on passenger fares for full-service airline passengers, and stochastic for all other passengers.

- For full-service airline passengers: for the final leg of their journey, all full-service airline leg fares are normalised from 0 (lowest) to 1 (highest), this value being labelled the normalised fare (f_N). This is then used as a weight for the soft cost value in the previous table, with the base cost scenario value at the lower end (i.e. assigned to the lowest weights) and the high cost scenario at the upper end (i.e. assigned to the highest weights). The delay is that (dynamically) estimated at the final destination.

- For all other airline types: for each passenger a soft cost is assigned, drawing on a Normal distribution, using the mean of the base and low cost scenario values as the mean of the Normal distribution, with a standard deviation set at one quarter of the difference between the base value and the low value. By the known properties of the Normal, this means that 95.5% of the sampled values will lie between the low and base cost scenario values. The delay is that (dynamically) estimated at the final destination.

Table 18 summarises the soft cost assignments by airline type.

Table 18. Summary of soft cost assignments by airline type

Aircraft operator type	Method type	Summary of method
FSC	Deterministic	f_N x high + [1-f_N] x base
REG, LCC, CHT	Stochastic	~N([base + low]/2,[base-low]/4)

For further derivation details and critiques of this methodology, please refer to [22] and [23]. Although the model estimates soft costs for each passenger dynamically, we can illustrate the magnitudes involved based on a soft cost applied equally to all passengers on a flight (taken from [23]). Although by two hours the per-minute soft cost has saturated, the net cost continues to increase with total length of delay. If we take a British Airways Boeing 747-400 on a return flight from London to New York, with an average four-class configuration and 80 per cent load factor, this would typically generate revenue of the order of €300k. The soft cost of a 5-hour delay (using the base cost scenario value in the table) is €94k — this equates approximately to the loss in revenue of one in three passengers taking their custom to another airline for one London–New York return trip.

3.3.1.15 APU fuel burn

APU fuel burn as shown below, with fit against √MTOW for other aircraft. APU is used up to 20 minutes during turnaround. Table 18 has the fuel consumption per aircraft type.

Table 19. APU fuel consumption per aircraft type

Aircraft	√MTOW	APU fuel kg/min
B733	7.77	1.98
B734	8.09	1.98
B735	7.46	1.98
B738	8.64	1.98
B752	10.38	1.98
B763	13.47	3.12
B744	19.82	3.66
A319	8.19	1.98
A320	8.60	1.98
A321	9.31	1.98
AT43	4.10	1.14
AT72	4.71	1.14
DH8D	5.40	1.14
E190	6.98	1.98
A332	15.18	3.66

3.3.1.16 Crew and maintenance costs

The top 15 of most used aircraft types are calibrated individually to calculate crew and maintenance costs (see the following three tables). For other aircraft types the model uses an interpolation based on √MTOW.

Table 20. Crew arrival delay costs

Aircraft	Cost per minute (€)	√MTOW
B733	8.9	7.77
B734	9.2	8.09
B735	8.4	7.46
B738	9.5	8.64
B752	9.9	10.38
B763	13.0	13.47
B744	17.5	19.82
A319	7.7	8.19
A320	8.2	8.60
A321	8.2	9.31
AT43	5.9	4.10
AT72	6.4	4.71
DH8D	6.4	5.40
E190	7.1	6.98
A332	13.8	15.18

Table 21. Maintenance at-gate costs

Aircraft	Cost per minute (€)	√MTOW
B733	0.50	7.77
B734	0.54	8.09
B735	0.46	7.46
B738	0.45	8.64
B752	0.59	10.38
B763	0.84	13.47
B744	1.17	19.82
A319	0.57	8.19
A320	0.52	8.60
A321	0.63	9.31

Aircraft	Cost per minute (€)	√MTOW
AT43	0.21	4.10
AT72	0.29	4.71
DH8D	0.29	5.40
E190	0.43	6.98
A332	0.89	15.18

Table 22. Maintenance en-route costs

Aircraft	Cost per minute (€)	√MTOW
B733	3.91	7.77
B734	4.21	8.09
B735	3.63	7.46
B738	3.39	8.64
B752	4.55	10.38
B763	6.22	13.47
B744	8.71	19.82
A319	4.34	8.19
A320	3.96	8.60
A321	4.75	9.31
AT43	1.71	4.10
AT72	2.31	4.71
DH8D	2.24	5.40
E190	3.34	6.98
A332	6.55	15.18

3.3.2 Door-to-door simulation

The door-to-door simulation model has two different data types that have to be calibrated: passenger archetypes (see Section 2.3.2) and airport archetypes (see next paragraph), in order to validate the model for 12SEP14. The model was calibrated previously as part of the Horizon 2020 DATASET2050 project.

The model includes 200 airports and has assessed granular information about the different internal process times (e.g. check-in, security) of the top ten airports (LFPG, EGLL, EDDF, EHAM, LFPO, EDDM, LEMD, LIRF, LEBL, EGKK). The rest of the airports have been aggregated. We have also plotted average walking distances within airports, therefore clustering them into two archetypes (large and small). These two types of airports are themselves clustered into two groups depending on their seasonality (i.e. if the traffic is higher in summer or winter due to some touristic location for example). These characteristics determine the passenger archetype dwell times at the airports.

4 Scenarios

4.1 Scenarios modelled in Vista

The scenarios potentially simulated by Vista were defined in D3.1. For this deliverable, we selected some of them to be run by the different blocks.

The selection of scenarios is based on the consultation, the workshop and the current models' capabilities.

We used five background scenarios: '**current**', '**L35**' (low economic growth and low technological development in 2035), '**H35**' (high economic growth and high technological development in 2035), '**L50**', and '**H50**' (*idem.* in 2050).

The background scenarios, as shown in Figure 62, allow us to analyse the evolution of the current system, with both low and high economic and technical development. From the feedback obtained during the realisation of the project, the low scenarios are considered more relevant and will be analysed in more detail, as explained below. On the other hand, the high baseline scenarios will allow us to push the system further from its current performance (higher demand and technological capabilities) enabling us to compare the metrics evolution and to present the capabilities of the Vista model to capture this higher disruptive case. Note that when referring to the 'baseline' without further information, we refer to the 'current' – 'low 2035' – 'low 2050' evolution.

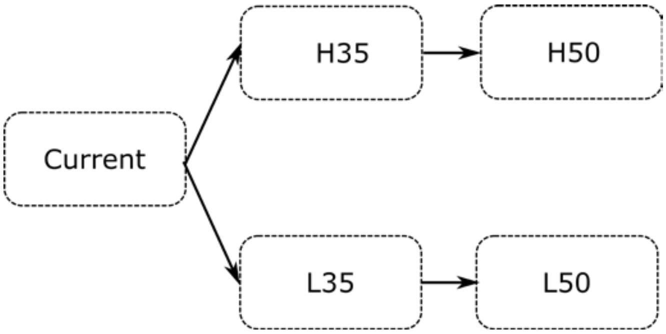

Figure 62: Five foreground factors modelled

Besides the modelling of the baseline system evolution, in Vista, five foreground factors have been modelled, as summarised in Table 23 and described in the rest of the section (see Section 4.3). The foreground factors have been grouped to create two possible foreground evolutions to be considered in the model: the supportive and non-supportive foreground factors evolution. Table 25 presents the values that the different foreground factors take when modelling supportive and non-supportive cases for 2035 and 2050. Note that supportive and non-supportive foreground scenarios are supportive or not in the sense that they benefit or penalise one or several stakeholders (e.g., in the supportive cases, the low fuel price is beneficial to the airlines but 'high' emissions scheme implementation should be good for the protection of the environment; while the opposite is true in the non-supportive case: very high fuel prices are detrimental for the airlines). However, such support might create unexpected interactions between different indicators that will be captured by the model.

As described previously, these foreground factors will be individually modelled only on the low 2035 and low 2050 background scenarios (baseline system evolution).

Table 25 shows all the scenarios (9) which have been simulated in the final version of the model. All of them have been run by all the layers, the output of one layer being the input of the next one (except for the final, tactical layer). The four last scenarios in the table are new and consist of different factors, sometimes opposed to each other, being implemented at the same time.

Table 23. Five foreground factors modelled

Foreground factors	Description
BEO1	Price of fuel (€/kg): 0.3 ('L'), 0.5 ('D', current), 1.0 ('H'), 2.0 ('HH'), 4.0 ('HHH')
BTS5	Implementation level of 4D trajectories: 'D' (current), 'L', 'M', 'H'
BTO4	Implementation level of passenger reaccommodation tools: • 'L' (current 'rule of thumb' operations); • 'M' (flexible reaccommodation for FSC; medium wait for passengers); • 'H' (fully flexible passenger reaccommodation).
ROR1	Implementation level of passenger provision schemes: • 'L' (current Regulation 261); • 'M' (enhanced Regulation 261 for passenger compensation); • 'H' (automatic compensation for entitled passengers).

Foreground factors	Description
ROR3	Implementation level of emissions scheme: • ‘D’ (current emissions scheme, low price, only CO_2 charged); • ‘CO2_H’ (high price, only CO_2 charged); • ‘CO2_HH’ (very high price, only CO_2 charged); • ‘CO2_NOx_H’ (high price, NO_x also charged); • ‘CO2_NOx_HH’ (very high price, NO_x also charged).
Key	‘D’ = default ‘H’ = high; ‘HH’ = very high; ‘HHH’ = extremely high ‘M’ = medium ‘L’ = low

Table 24: Foreground factor values: default (baseline), supportive & non-supportive scenarios

	Current	Supportive 2035	Supportive 2050	Non-supportive 2035	Non-supportive 2050
BEO1 – fuel price	D	L	L	HH	HH
BTS5 – 4D trajectories	D	H	H	L	L
BTO4 – Pax management tools	L	H	H	L	L
ROR1 – Pax provision	L	H	H	L	L
ROR3 – Emission scheme	D	CO2_NOx_H	CO2_NOx_HH	CO2_H	CO2_HH

Table 25. Scenarios simulated by Vista

Scenario	Description
Current	Current situation (SEP14)
L35 baseline	Baseline environment in 2035 with slow economic growth and slow technological development
H35 baseline	Baseline environment in 2035 with fast economic growth and fast technological development
L50 baseline	Baseline environment in 2050 with slow economic growth and slow technological development
H50 baseline	Baseline environment in 2050 with fast economic growth and fast technological development
Non-supportive 2035	Poor emphasis on environmental and passenger protection, very high price of fuel in L35
Supportive 2035	Strong emphasis on environmental and passenger protection, low price of fuel in L35
Non-supportive 2050	Poor emphasis on environmental and passenger protection, very high price of fuel in L50
Supportive 2050	Strong emphasis on environmental and passenger protection, low price of fuel in L50

4.2 SESAR-related business factors (BTS)

As explained in D2.1, the factors have been bundled together and some of them are easier than others to implement in the model. In particular, all the 'BTS' factors, which are linked to operational improvements as envisioned by SESAR, have been taken into account using the available information. Indeed, as explained in D2.1, we used the high-level targets or actual performance assessments on KPIs from SESAR for subpackages to estimate their impact. For instance, the SESAR subpackage 'Enhanced Runway Throughput' (SPC01.03) corresponds to the Vista factor BTS3 and had a performance assessment, showing a potential benefit of +5.62% for airports capacity (for time-based operations).

The different targets and performance assessments are thus used to estimate what would be the future benefits delivered. Doing so, we made some assumptions on the future for the different time-horizons. We considered the previous 'Step1/2/3' schemes designed by SESAR, and their targets. Since not all subpackages have targets for the three steps, we used some interpolation and extrapolation to compute the missing data:

1. We compiled for Step1 the most up-to-date information, i.e. the results of the assessment reports where there were some (where not, we used SESAR targets);
2. For each KPI, we computed the overall improvements for Step1 and Step2. For instance, airport capacity should be increased by 8% with Step1 and a further 5% with Step2;
3. We computed for Step1 the contribution of each subpackage to the overall target. For instance, SPC01.03 contributed up to 78% to the overall target of 8%;
4. For the KPIs which are completely missing in Step2 (predictability and en-route capacity), we used the proportions found in Step1 and applied them to the overall Step2 target. For instance, SPC03.01 was expected to trigger a 3.05% improvement of en-route capacity;
5. We performed the same interpolation with the Step3 high-level targets.

The results of this procedure are shown in Annex 8.

We then made the link between Step1/2/3 and the different scenarios built in Vista, as explained in D3.1. We considered indeed that 'low' values ('L') for the factors correspond to time-based operations (Step1), 'medium' ('M') to trajectory-based operations (Step2) and 'high' ('H') to performance-based operations (Step3). Based on the background scenarios we define in D3.1, we thus have the correspondence shown in the next table.

Table 26. Correspondence between background scenarios and operation types

Background scenario	Operation type
Current	No improvement
L35	Time-based
M35	Time-based
H35	Trajectory-based
L50	Trajectory-based
M50	Trajectory-based
H50	Performance-based

The various blocks use this information differently by changing parameters in their respective models based on their values, as described in D3.1, or more explicitly by implementing a different mechanism based on the underlying operational improvements.

4.3 Non-SESAR factors

4.3.1 BEO1 – price of fuel

To begin with, we use the price of fuel (BEO1) as a foreground factor. In Table 27 we show the values of the price of fuel for the different settings of BEO1.

Table 27: Price of fuel values

BEO1 value	L	D	H	HH	HHH
Price of fuel	0.3	0.5	1	2	4

4.3.2 BED1 – economic development

For BED1, we estimated the economic growth of the countries and its impact on the demand. For this, we used the projection from STATFOR [10] in terms of GDP growth, with the different scenarios they define (fragmenting world, regulated growth and global growth, respectively applied in the low, medium and high scenarios). We applied different growths for different countries, using averages when data were missing, as shown in the following table. We then used GDP per inhabitant to estimate the changes of income of the passengers in the economic model.

Table 28. GDP growth in the different scenarios

BED1 value	STATFOR scenario used	Background scenario impacted	Overall growth
D	N/A	Current	0%
LL	Fragmented world	L35	25%
LM	Happy localism	M35	40%
LH	Global growth	H35	48%
HL	Fragmented world	L50	48%
HM	Happy localism	M50	77%
HH	Global growth	H50	90%

4.3.3 ROR3 – emissions scheme

ROR3 assesses potential changes in the European Emissions Trading System (ETS), and worldwide. The current scheme works as a system of allowances, which can be traded on a secondary market. The airlines also have a certain amount of free allowances.

For the current price of allowances, we use [15], in which a temporal evolution is provided; airlines had 82% of free allowances. Given that the price of allowances was 5 euros per ton 2014, we used an effective price of 5 x 0.20 = 0.12 euros per ton for our current scenario.

For future scenarios, we used several projections. Reference [16] gives some long-term values and mentions others used by the French government. These have similar values. However, a more recent publication [17] predicts lower prices for the future. In fact, in the highest scenario of this publication, the prices are roughly at the level of the lowest scenario for [16]. As a consequence, we used these values, which are equal to 50 and 100 euros per ton, respectively. Moreover, we decided to consider that airlines would pay for a higher share of their allowances in the long-term future. We thus chose to consider that they would pay for the same amount in 2035 (20%), and 50% in 2050. This gives as effective prices 12 and 50 euros per ton. Table 29 shows a summary of this procedure.

Table 29. ROR3 values explanation

ROR3 value	Allowance price (€/kg)	Percentage of non-free allowances	Effective price (€/kg)	NO_x charged	H_2O and contrails charged	Used in final model
D	0.005	20%	0.0012	No	No	Yes
CO2_H	0.05	20%	0.012	No	No	Yes

ROR3 value	Allowance price (€/kg)	Percentage of non-free allowances	Effective price (€/kg)	NO_x charged	H_2O and contrails charged	Used in final model
CO2_HH	0.1	50%	0.05	No	No	Yes
CO2_NOx_H	0.05	20%	0.012	Yes	No	Yes
CO2_NOx_HH	0.1	50%	0.05	Yes	No	Yes
All_H	0.05	20%	0.012	Yes	Yes	No
All_HH	0.2	50%	0.05	Yes	Yes	No

4.3.4 ROR1 – passenger provision

Regulation 261/2004 establishes the minimum rights for passengers when they are denied boarding, their flight is cancelled or delayed. This includes right of care and the right to compensation [20]. The European Commission will adopt interpretative guidelines in order to provide guidance to the citizens and the airlines on current passengers' rights. In February 2014, the proposed strengthening (*inter alia*) of air passenger rights passed its first reading in the European Parliament [21] in the following directions in particular: right to care, re-routing, connecting flights. Moreover, the European Parliament's proposals also go further than those proposed by the Commission in strengthening these air passenger rights (*ibid*.). Other possible evolutions of passenger provision regulations include:

- passengers entitled to compensation being automatically compensated;
- load factors maintained significantly below 100% on key/connecting/trunk routes to reserve some capacity for rebooking passengers who miss flights/connections - a 'social' capacity and resilience provision supporting Flightpath 2050 ambitions through new regulatory paradigms;
- enhanced identification of primary delay reasons to adjust airline liability.

Table 30 through Table 33 summarises the Vista model assumptions.

Table 30. Duty of care costs as a function of aircraft operator and ticket type

AO (& ticket) type	Cost applied
FSC (**premium**)	average of high and base
FSC (**standard**)	base
REG, LCC, CHT	average of low and base

Table 31. Cost due to duty of care, euros per passenger by minutes of delay

	ROR1 L			ROR1 M & H		
Minutes of delay	**high**	**base**	**low**	**high**	**base**	**low**
0 - 90	0	0	0	0	0	0
90 - 120	2.2	1.8	0.0	10.2	8.4	5.1
120 - 180	10.2	8.4	5.1	10.2	8.4	5.1
180 - 300	25	21	13	25	21	13
300 - 480	28	23	14	28	23	14
480 - **Overnight**	109	90	55	109	90	55

Table 32. Reg261 compensation

	Flight type	**Delay (in minutes)**		
		120 < 180	**180 < 240**	**>240**
ROR1 L	Short haul (<1,500 km)	0	€250	€250
	Medium haul (>1,500 km & < 3,500 km)	0	€400	€400
	Long haul (>3,500 km)	0	€300	€600
ROR1 M & H	Short haul (<1,500 km)	125€	€250	€250
	Medium haul (>1,500 km & < 3,500 km)	125€	€400	€400
	Long haul (>3,500 km)	125€	€300	€600

Table 33. Passenger claiming compensation ratios in ROR1

	% of passengers entitled claiming compensation
ROR1 L	11%
ROR1 M	60%
ROR1 H	100%

4.3.5 BTO4 – passenger management tool

Improved passenger reaccommodation tools will affect how passengers are reaccommodated according to a better understanding of the cost of delay incurred by airlines. As more factors are considered, actual (aircraft and passenger) waiting times and missed connections could be either increased or decreased, as not simply the delays but their *costs* will drive these processes.

Currently, stranded passengers in the tactical layer are reaccommodated based on their ticket category and the airline archetype. In general, full-service carriers will only reaccommodate premium passengers onto other alliance flights, all other passengers are reaccommodated within the alliance unless the delay is greater than five hours. In the baseline strategic layer, basic rules of thumb are modelled. However, in future work, delays of some key flights could be increased to wait for premium passengers, for example.

4.4 Strategic layer

The economic model takes into account a number of factors that impact some values in the simulation.

First, it uses all BTS factors, i.e. the SESAR improvements, using the high-level targets as described in Section 4.2. For each KPI, the value is then converted into a modification of a parameter of the model. Table 34 summarises the effects of KPIs on the model.

Table 34. Impact of SESAR KPIs on inner parameters of the economic model

KPI	KPI description	Model parameter affected	Parameter Description
APT_CAP	Airport capacity increase	$\mathrm{C} \rightarrow (1+\mathrm{x})\mathrm{C}$	Empirical capacity coefficient
ASP_CAP_ER, ASP_CAP_TMA	Airspace capacity increase, en-route and TMA	$C_{max} \rightarrow (1+\mathrm{x})C_{max}$	ANSP pre-factor in delay generation function
PRE	Predictability increase in airport departures	$\sigma_{oo} \rightarrow (1+\mathrm{x})\,\sigma_{oo}$ $\sigma_s \rightarrow (1+\mathrm{x})\,\sigma_s$	Airport slope and ordinate at origin for standard deviation for delay versus mean delay at airport
ENV	Geometrical inefficiency decrease of trajectory	$b \rightarrow (1+\mathrm{x})b$ $\mathrm{c} \rightarrow (1+2\mathrm{x})c$	Linear and quadratic coefficients for cost of fuel
EFF_TECH	ATCO inefficiency decrease in terms of cost per ATCO	$\omega \rightarrow (1+\mathrm{x})\omega$	Cost per unit of empirical capacity
EFF_ATCO, EFF	ATCO efficiency increase in terms of number of ATCO per kilometer flown	$\omega \rightarrow \omega/(1+\mathrm{x})$	Cost per unit of empirical capacity
PUN	Punctuality increase in terms of number of flights below 3 minutes of delay	$\delta t_0 \rightarrow (1-\mathrm{x})\delta t_0$	Average airport delay generated at zero traffic

In terms of foreground factors, we consider BEO1 (the price of fuel), BTS5 (4D trajectory operations) and ROR3 (emissions scheme). The latter is directly controlled by the parameters p_e and $e_{NOx,}$ in the economic model, which enter the operational cost function of the airline. Note that compared to the initial list of foreground factors applicable to the economic model, only BEO2 (different CRCO charges schemes) and RAD2 (regional airport development) have not been considered in this final version.

In Table 35 we summarise how non-SESAR factors impact the model.

Table 35. Effect of other factors on inner parameters of the model

Factor	Factor description	Model parameter affected	Parameter description
BEO1	Fuel price	p_f	Fuel price in airline cost function
BED1	Economic development	i	Income of passengers
ROR3	Emission scheme	$p_e, e_{NOx}, e_{H2O}, e_{con}$	Emission price for each type of emission

4.5 Pre-tactical layer

The pre-tactical layer takes as input the outcome of the strategic layer. Therefore, factors that affect the demand (both for flights and passengers) and the capacity of the ANSPs, are already considered as they affect the schedules for which flight plans will be generated, the flow of passengers and the capacities of the ANSPs.

Then other factors will impact directly the pre-tactical layer models. Table 36 shows the impact of the corresponding SESAR KPI (Geometrical inefficiency decreases of trajectory) on the pre-tactical layer blocks and Table 37 summarises the different factors that are explicitly considered in the pre-tactical layer, how they are considered and their expected impact on the model.

Table 36. Impact of SESAR KPIs on the pre-tactical layer

KPI	KPI description	Block affected	Summary description	Expected impact
ENV	Geometrical inefficiency decreases of trajectory	Flight plan generator – route generation	The flight plan distance of the pool of routes is reduced by this environment improvement; each segment within each ANSP of the routes are reduced by this factor too	Lower fuel consumption and flight time needed per origin-destination pair

Table 37. Summary of factors modelled in pre-tactical layer

Factor	Block affected	Summary description	Expected effect
BEO1 – Fuel price	Flight plan generator – flight plan generation	Fuel prices have an impact on cost of the different flight plan options for each flight	Higher fuel values will lead to the preference of more direct and shorter routes even if the total CRCO cost is increased Low price of fuel will have the opposite impact, encouraging airlines to select longer routes that use cheaper airspaces when possible Note that the current value of fuel is already relatively low, for this reason it is expected that decreasing even further the cost of fuel might not have a high impact on the system
ROR3 – Emissions scheme	Flight plan generator – flight plan generation	The emissions scheme has an impact on the cost of CO_2 and hence on the cost of the different flight plan options.	As with fuel (BEO1), higher emissions scheme value might lead to a preference for routes which are more direct and use less fuel.

4.6 Tactical layer

A set of five factors has been modelled in the tactical layer. The factors affect several aspects of the tactical model: airports, airspace and passenger management, which, in turn, demonstrates the flexibility of the tactical layer. In addition, when not explicitly modelled, all mechanisms also impact the tactical layer, any change in the strategic or pre-tactical layer ultimately has an effect on the tactical layer. Note that the precise values of the parameters are extracted from the expected effects described in Section 3 of D2.1.

As the definition of each factor is rather conceptual, and can be interpreted in many different ways, only the impacts of changing the model parameters are assessed. The tactical layer is a 'what-if' assessment tool in the sense it is subjugated to the parameters. More interesting is the intricate effect of having more than a single factor simultaneously. In a first phase, each factor is simulated in isolation. This produces an isolated factor analysis which can be compared with the baseline results for calibration and impact assessment. Once the baseline and factors are calibrated, all combinations of two and three factors are considered. The interactions are far more complex when considering multiple factors simultaneously, and calibration is no longer possible, but the results allow us to explore synergies and discordant effects across the factors.

4.6.1 Fuel prices (BEO1)

Volatility of fuel prices drives the tactical behaviour of the airspace users. The aircraft cost index is directly related to the actual cost of fuel, so in a scenario of higher fuel prices it is most likely that airline operators are reluctant to recover delays. The opposite is also true, in a low fuel price scenario, more delay is likely to be recovered. The expected effect on the tactical model is therefore focused on increasing/decreasing gate-to-gate predictability by allowing the pilots to recover/allow more delay.

In the baseline tactical model most airline operators use either a general rule of thumb or cost indexing to recover delay: either way, the amount of delay is limited to a % of the nominal route length and a residual delay that is not worth recovering due to the low impact on total (delay) costs. This factor is modelled through two parameters, the maximum % of delay to be recovered as a function of the nominal route length and the amount of residual delay. The implementation is summarised in the following table.

Table 38. BEO1 - Fuel prices

BEO1 - Fuel prices	HHH	Eight times current price	• up to 0.5% delay recovered • 30 minutes of residual delay
	HH	Four times current price	• up to 2% delay recovered • 20 minutes of residual delay
	H	Double current price	• up to 3% delay recovered • 15 minutes of residual delay
	D	Current market price	• up to 5% delay recovered • 10 minutes of residual delay
	L	Half current price	• up to 7% delay recovered • 5 minutes of residual delay

4.6.2 4D trajectory operations (BTS5)

As described in Section 3 of D2.1, the factor BTS5 includes the OFAs OFA03.01.03 (Free Routing) and OFA03.01.04 (Business and Mission Trajectory) which calls for an increased integration between the NOP and the Airport Operations Plan by sharing operation planning and real-time data, and predicting traffic complexity and overloads with planned trajectory and network information. This will allow application of mitigation strategies at local and network levels.

In the tactical model, these operational and technological changes affect both the airport and airspace capacity as well as producing an increased predictability for the gate-to-gate phase. Each of these elements is modelled differently in the tactical layer and therefore is affected by a different set of parameters. In the case of airport capacity, an average runway occupancy time (ROT) is used. Increase in airport capacity is simply modelled by aligning a scaling factor to the ROT. (Note that the tactical layer does not address how traffic flows would change if local capacity at airports is increased – this is done in Vista at the strategic layer.)

En-route capacity is modelled by reproducing the non-reactionary distribution of delay, in this case using the conditional probability of the chance of a given delay being ATFM. Final en-route predictability is modelled by adding random noise to the nominal duration of a flight, in this context an increase in predictability is simply a reduction on the standard deviation of the random noise applied.

4.6.3 Passenger management tools (BTO4)

See Section 4.3.5 for details of this implementation, which is summarised in the following table. In each case, aircraft only wait if their load factor is <80%, to avoid already-boarded passenger costs downstream.

Table 39. BTO4 - Passenger management tools

BTO4 - Passenger management tools	L	Current rule of thumb operations	Recover up to 10 minutes of residual delay
	M	Flexible accommodation for FSC; medium wait for passengers	Passengers reaccommodated on any flight to destination for FSC tickets Medium wait for passenger rules: 30 minutes for outbound international connecting passengers
	H	Fully flexible passenger accommodation	Passenger reaccommodated on any flight to destination High wait for passenger rules: 60 minutes for outbound international connecting passengers

4.6.4 Passenger provision schemes (ROR1)

The Vista tactical layer is affected in two ways by modifications to Regulation 261: (i) by considering a different cost schema for airlines due to passenger regulations, e.g. enhanced passenger compensation, and, (ii) by implementing automatic compensation for entitled passengers. Since the current claims for compensation are far from 100%, it is expected that airlines might behave differently by estimating the cost and hence this may have an impact on delay recovery and wait-for-passenger rules. The implementation is summarised in Table 40.

Table 40. ROR1 - Passenger provision schemes

ROR1 - Passenger provision schemes	L	Current R261	
	M	Enhanced R261 for passenger compensation	• Increased duty of care in the range 90-120 to 90-180 minutes • 60% claim rate
	H	Automatic compensation for entitled passengers	• Increased duty of care (as above) • 100% claim rate

4.6.5 Emissions scheme (ROR3)

A modification of the emissions scheme will have a direct operating cost for each flight plan. This is an effect in the pre-tactical phase, when flight plans are calculated. However, similar to the BEO2 mechanism, the effects are propagated to the tactical layer since the impact on probabilities of selecting different flight plan options changes for cost-sensitive selections. Some airlines may adopt an environmentally friendly policy to compensate for higher emissions scheme costs, however, the scenario considered in Vista is an increase in allowances for CO_2 and NO_X. There is, however, a direct effect on the tactical layer metrics, since the operating costs of each flight execution is affected by emission charges.

5 Results

In the following sections, we describe the results for each layer, firstly by the low and high baselines, then for the supportive and non-supportive scenarios (NB. comparing with the *low* baselines only, to reduce the volume of the outputs), followed by summaries of the key findings.

5.1 Strategic layer

The scenarios tested in the final version for the strategic layer aim to estimate the joint impacts of different, sometimes opposing, factors.

Regarding metrics, we used several ones per agent, as listed below.

- Airport:
  - capacity;
  - (departure) delay (per flight);
  - total costs;
  - profit;
  - revenues;
  - traffic (for 24 hours).
- Airline:
  - (departure) delay (per flight);
  - ATFM delay (per flight);
  - ATC cost (per flight);
  - airport cost (landing/terminal fees, per flight);
  - cost of delay (per flight);
  - emission costs (per flight);
  - fuel cost (per flight);
  - crew cost (per flight);
  - maintenance cost (per flight);
  - cost of uncertainty (per flight);

  - cost of uncertainty per minute (per flight);
  - emissions in CO_2 equivalents;
  - ticket price (per OD pair);
  - size of aircraft (average number of seats per flight);
  - fuel consumption (per OD pair);
  - uncertainty (per flight);
  - number of passengers (per OD pair);
  - total operational cost.

- ANSP:
  - capacity;
  - delay;
  - expenses;
  - maximum capacity;
  - revenues;
  - traffic;
  - unit rate;
  - spare capacity (fraction of capacity that is not used with respect to traffic);
  - slack capacity (fraction of capacity left before reaching theoretical maximum capacity allowed by technology).

- Pax (all per itinerary):
  - airport delay;
  - ATFM delay;
  - ticket price;
  - income;
  - fare to income ratio (per week);
  - frequency;
  - volume.

Due to the number of metrics, displaying and analysing all of them is too cumbersome. In the following, we thus highlight the main, latest results. All the results can be found in Annex 9 (including results previously reported in D5.1).

Note that in the simulations, a typical number of runs is around 200 to ensure convergence. In terms of number of agents, the simulations feature:

- 823 airport agents;
- 326 airline agents;
- 15 209 flight agents;
- 31 430 passenger agents;
- 88 ANSP agents (but only the ECAC ones are active).

5.1.1 Overview

We begin this analysis by showing two maps in Figure 63 and Figure 64, illustrating briefly the scope of the model and what is going on in the system. In this map, we drew (and superimposed) a line for each flight. We coloured the line in red if we expected this flight to be late, otherwise the line is white. The line is completely white if the delay is under 5 minutes, and completely red if the delay is above 15 minutes.

Figure 63. Map of the state of the system in the current scenario

Comparing both maps, set in the current scenario and in the H35 baseline, respectively, it is striking to see how traffic increases in the second case, and how flights are much more likely to be delayed. Interestingly, this phenomenon is not uniform. Traffic seems to increase a lot more in eastern Europe and the Middle East, than in southern Europe, for instance. Traffic coming from and going to eastern Europe is a lot more delayed than other flights.

As shown in Annex 9, the situation is different again for other baselines (the most stressed one being H50) and the supportive and non-supportive scenarios. In the following, we study more in detail how these situations differ.

Figure 64. Map of the state of the system in the H35 baseline scenario

5.1.2 Baseline evolution

We start by studying the evolution of the metrics for the baselines.

5.1.2.1 Low baseline

In Figure 65, we show a plot comparing the different metrics measured for airports. In this figure, capacities are not given in number of movements, but are scaled with the number of passengers. The traffic is also given in number of passengers. The delay is in minutes.

All the metrics follow the same pattern, smaller values for the current scenario, then higher for 2035. This is mainly the effect of traffic. Indeed, the traffic increases substantially between the current situation and 2035, and rather less between 2035 and 2050. Interestingly, however, one can note that the capacity of the airports increases a lot as well, due to technological improvements. This could translate into decreasing delay at the airport, which apparently is not the case in these scenarios: the gains in capacity cannot offset the loss in terms of traffic.

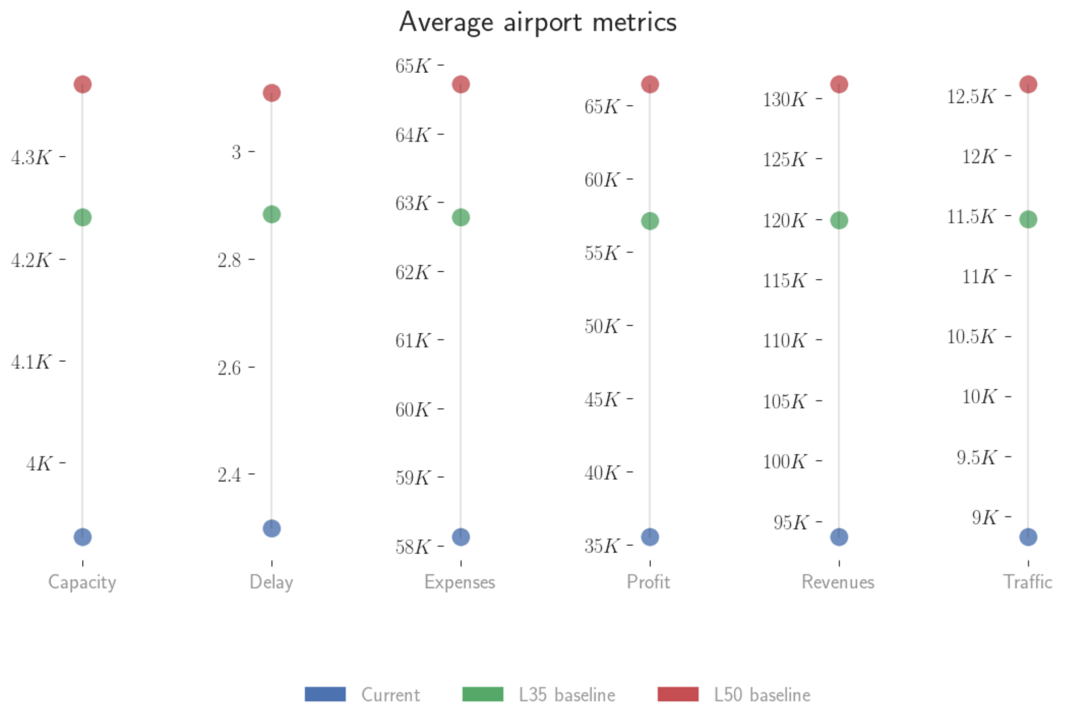

Figure 65. Average metrics for airports

In Figure 66, we show a similar plot for the airlines. In this figure, all delays and uncertainty are given in minutes. “C.p.f” stands for “cost per flight”, “Unc.” for “uncertainty”, “p.m.” for “per minute”. The emissions are given in CO_2 equivalents (“eq. CO2”), per flight. “Fuel Con.” is the fuel consumption and “Tot. Op. C.p.f” is the total operational cost per flight, i.e. the sum of the other costs in this graph.

This plot confirms the comment regarding airport delay, but interestingly shows that ATFM delays show another pattern. Indeed, starting from 2 minutes, ATFM delay jumps to almost 3 minutes in 2035, before returning abruptly to little more than 2 minutes in 2050. The reason is that the traffic for 2035 is significantly higher than current traffic but stagnates (relatively) afterwards. On the other hand, technology continues to improve, and, in particular, the maximum capacities of ANSPs, which increase. As a result, some ANSPs hit their maximum capacity in 2035, which creates a lot of delay, but manage to return under this threshold (because the threshold has increased) in 2050. Note, however, how the ATC cost per flight keeps diminishing throughout the period, due to sizable increases in ANSP efficiency and thus decreased unit rates.

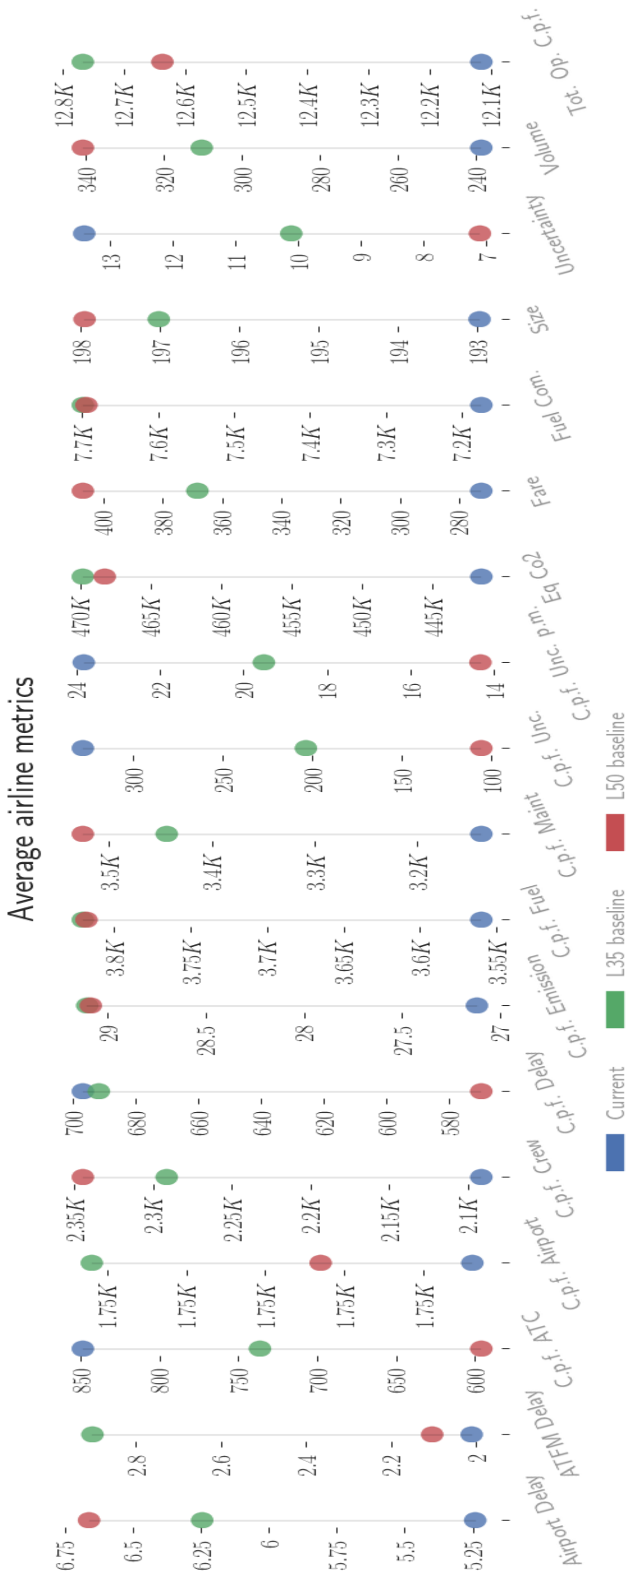

Figure 66. Average metrics for airlines

For this reason, the cost of delay shows a very moderate decrease in 2035, which is due *in fine* to the decrease in uncertainty, counter-balancing the increase in airport and ATFM delay. The decrease in uncertainty is indeed significant, as shown on the right of the figure, from a standard deviation of almost 14 minutes, down to 7 minutes in 2050. This is due to technological improvements, which are apparently more than enough to offset the increase in uncertainty due to the increase in mean delay.

The **cost of uncertainty** is estimated by calculating the difference of the cost of delay in two situations:

- the normal situation whereby the distribution of delay is driven by the airports;
- using a distribution that has the same mean as in the previous case, but with a vanishing (zero) variance.

The second cost coincides exactly with the cost of the expected (mean) delay, whereas the first is the expected cost of delay. The difference between both costs can be viewed as the **contribution of the standard deviation to the cost of delay**. This is probably an underestimation of the full cost of uncertainty, since an actual vanishing variance would trigger additional gains, such as a reduction of the mean itself due to stakeholders adapting their behaviours (e.g. airlines reducing buffers).

As shown in the figure, the cost of uncertainty represents nearly 50% of the cost of delay in the current scenario, and thus a reduction of the uncertainty leads to a significant drop in the cost of delay. While the cost of uncertainty decreases naturally with uncertainty, it is also interesting to see that the cost of uncertainty per minute decreases substantially too, by approximately 40%. Note also that the contribution of the cost of uncertainty to the cost of delay decreases from 50% to less than 20%, in 2050.

The fuel consumption follows an interesting pattern also, for different reasons. Indeed, the airlines tend to increase more their capacities on connections with heavy aircraft, shown by the increase in the size of the aircraft. These connections seem to be the most profitable for the airline, especially because the price of fuel is kept constant here. The consequence is an increase in the average fuel consumption for the airlines, as well as an increase in the cost of fuel and the cost of emissions (which is very small compared to other costs in these scenarios anyway). The gain in efficiency in terms of fuel consumption between 2035 and 2050 seem to be enough to offset the increase in aircraft size.

Note also how the total operational cost behaves as a result of this: the cost increases substantially in 2035, mainly driven by the increase in the long-haul flights and the increased delays. It decreases in 2050, driven by gains in efficiency and in uncertainty, to stabilise at a value significantly higher than the initial (current) one.

Figure 67. Statistical tests for airline metrics

When dealing with statistical data, it is also important to remember that similar but different averages can be the mere result of 'noise' (random effects). It is thus important to perform statistical tests. To do this, we used permutation tests, which are very robust and non-parametric. We used two-sample tests, i.e. we compare averages pairwise, one scenario against another. We used the Bonferroni correction to eliminate biases arising from multiple tests. The Bonferroni connection is the most conservative one for multiple test analysis, it comprises dividing the p-value (1% here) by the number of tests (three for instance for three scenarios, six for four scenarios, etc.). Some test results are shown in Figure 67 for the airline metrics. In this figure, each point represents a test, the null hypothesis for which is rejected (i.e. the averages *are* statistically different) if green, and red otherwise (i.e. the averages do not differ significantly). The tests are performed between the two scenarios on the abscissa and ordinate. "S0" represents the current scenario, "S1" represents the L35 and L50 baseline. For instance, there is no statistically significant difference between averages between the current scenario and the L50 baseline. The full tests for all the metrics computed in this deliverable can be found in Annex 9.

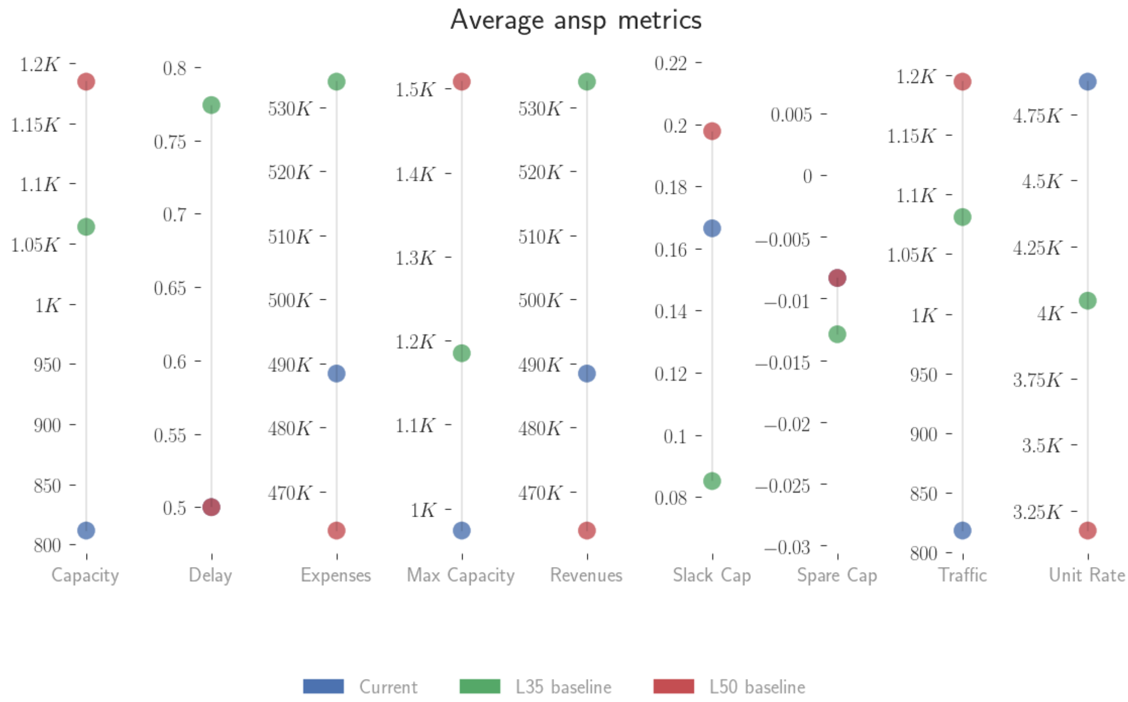

Figure 68. Average metrics for ANSPs

Figure 68 shows the point of view of the ANSPs. In this figure, capacities are given in number of flights and delays are in minutes. Spare and slack capacities are given in terms of fractions. As highlighted before, it is clear that the technological gains leading to higher maximum capacity are not fast enough to keep up with the increase of traffic in 2035: the traffic grows by around 35%, whereas the maximum capacities grow by only around 20%. As a result, the delay increases, the spare capacity is negative (-1.3%: meaning roughly that ANSPs would need 1.3% more capacity to attain their delay targets) and their slack capacity is very small on average (meaning that ANSPs cannot expand anymore with this level of technology). Thanks to the large increase in the maximum capacity between 2035 and 2050, all of the ANSPs attain their delay targets. Note that the delays presented are probably underestimated, as pointed out by Belgocontrol, in particular. The main reason for this is that the ANSPs in reality cannot expand and decrease their capacity at will (because of hiring, training, etc.), and that the traffic predictions they have access to are not sufficient. Also, note that the unit rate is in cents of euros.

We show in Figure 69 the passenger perspective. In this figure, delays are in minutes, “Fare” represents the ticket price for the whole journey (in constant euros), and “fare to income” is the ratio between the fare and passenger income (per week). The frequency here is the minimum of the frequencies on each leg of an itinerary. For a given OD, the average frequency is weighted by the number of passengers travelling on each itinerary. The “volume” is given as passengers per itinerary, per day.

The increase in the average income (completely exogeneous to the model) is the main driver for the overall increase in traffic. It is interesting to note that even though the ticket price increases, it increases less quickly than the income, as shown by the fare-to-ratio metrics.

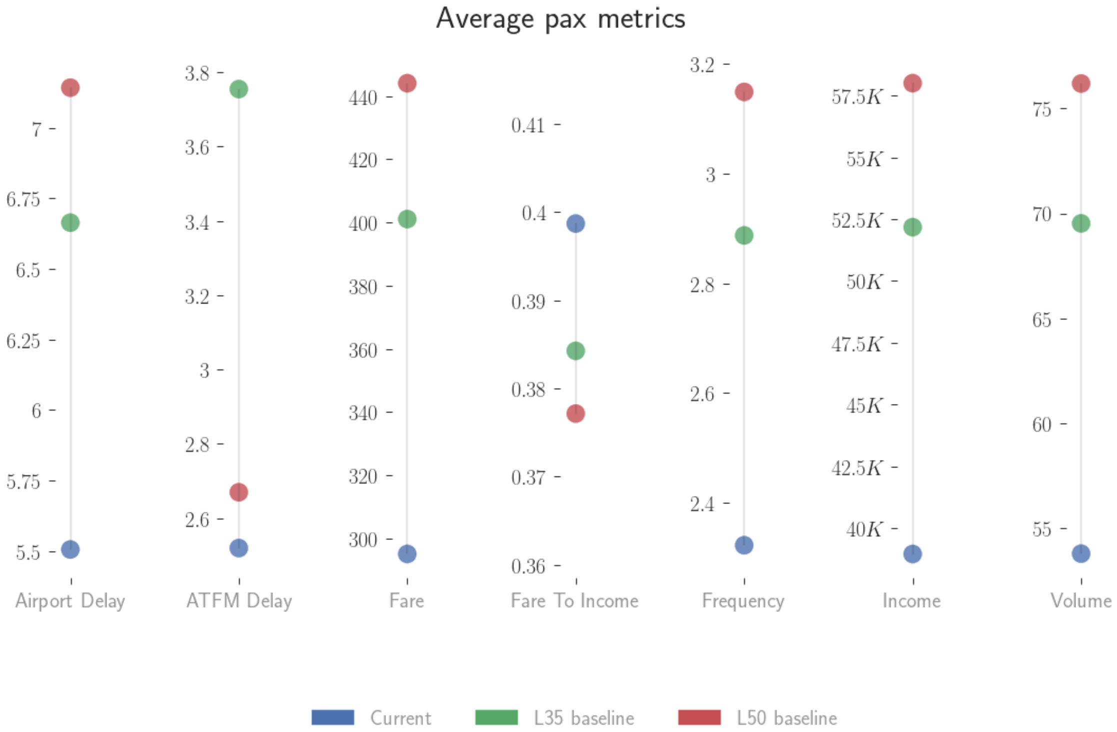

Figure 69. Average metrics for passengers

5.1.2.2 High baseline

In order to see the impact of the economic development and technological advancements, we now move to the baseline comparisons. Instead of showing the same kind of graphs, we use linear temporal plots for a few metrics to highlight the main differences between the two baselines ('low' and 'high').

We first illustrate in Figure 70 what is happening with the ANSPs and their capacities. On the upper left part of the figure, we show for the two baselines the capacity, the traffic, and the maximum capacity of the ANSPs on average. As explained, the ANSPs start with a maximum capacity that is 20% more than current capacity. In 2035, this 20% percent has reduced a little for the low baseline, which translates into a moderate delay (see upper right part of the figure). The traffic is only slightly higher than the capacity, which is normal since the target delay is non-null for ANSPs. Afterwards, the gap enlarges substantially and the delay falls again. For the high baseline, the gap is really small in 2035, with traffic significantly clearly higher than the capacity. This creates a sizeable delay, more than twice the delay in 2035 and more than three times the delay in the initial state. The system is slightly relaxed in 2050, but with a final delay which is significantly higher than the target (more than double). Note that these average values hide some very different situations for different ANSPs. Some are far from their maximum capacity, whereas others are very close to it.

In the lower part we show the evolution of the unit rates. In 2035 the unit rates drop faster for the high scenario, which is a direct consequence of an increase in traffic (ANSPs are profit-neutral, which means that unit rates decrease with traffic if capacity is fixed, and increase with capacity if traffic is fixed). Interestingly, both average unit rates end up approximately at the same level in 2050, which means that the efficiency in the H50 offsets exactly the traffic with respect to L50.

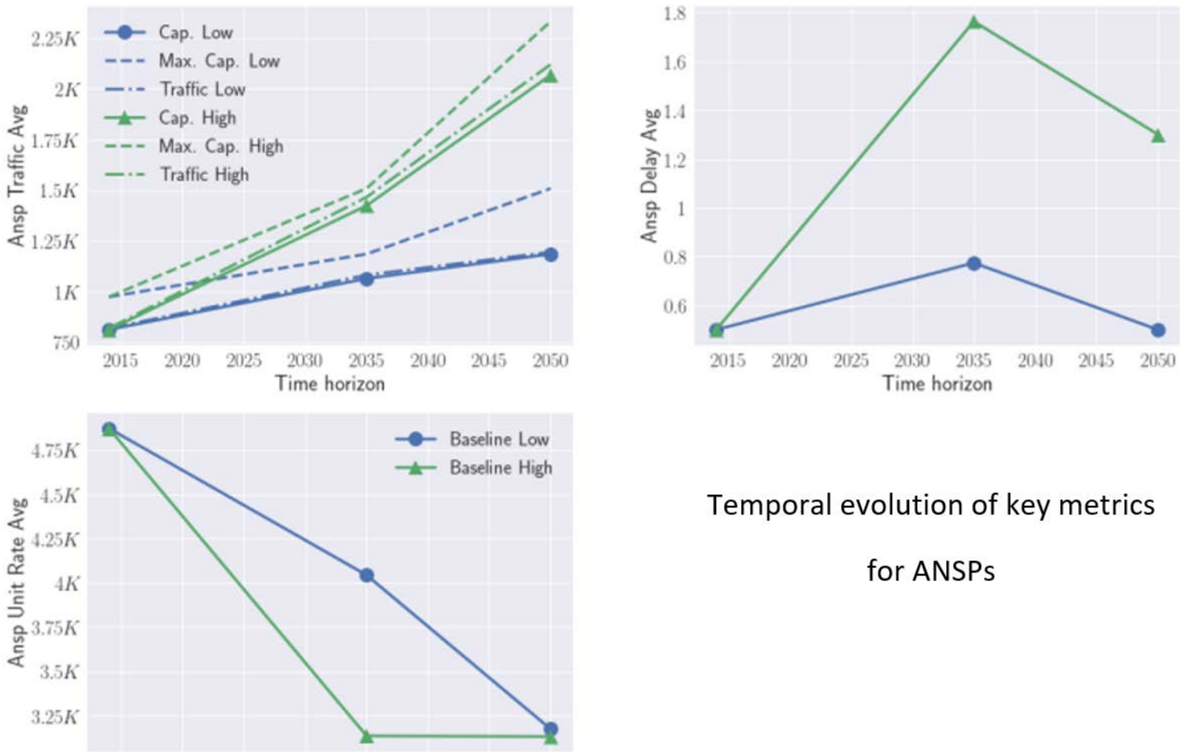

Figure 70. Temporal evolution of key metrics for ANSPs

In Figure 71 we continue by showing several key metrics for airlines. First, we see that the two kinds of delay are behaving differently (upper left plot). As seen above, ATFM delay on average tends to drop in 2050 for both baselines, whereas airport delay increases monotonically. In fact, the increase in delay is bigger in the second period than in the first, reaching more than twice the initial delay. This means that even with high technological developments, the increase of traffic is too high for the airports to cope with (unless they expand their infrastructure, which is not the case in the model). Technological gains, on the other end, seem to act faster for ANSPs and allow them to decrease their delay at least in the second period. For these reasons and others, the operational costs are increase substantially in the second period in the high scenario, whereas they drop a little in the low scenario (see upper right panel).

It is interesting to plot side by side the evolution of the fare and the evolution of the profit for airlines (second row in the figure). Fares increase quite a lot (but slower than income, see bottom left plot), but profits also, in a similar fashion. This points to the fact that competition (between airlines) in the model might not be strong enough. Indeed, in a perfect competition state, profits should tend to zero, since the fare should approach the marginal cost per passenger.

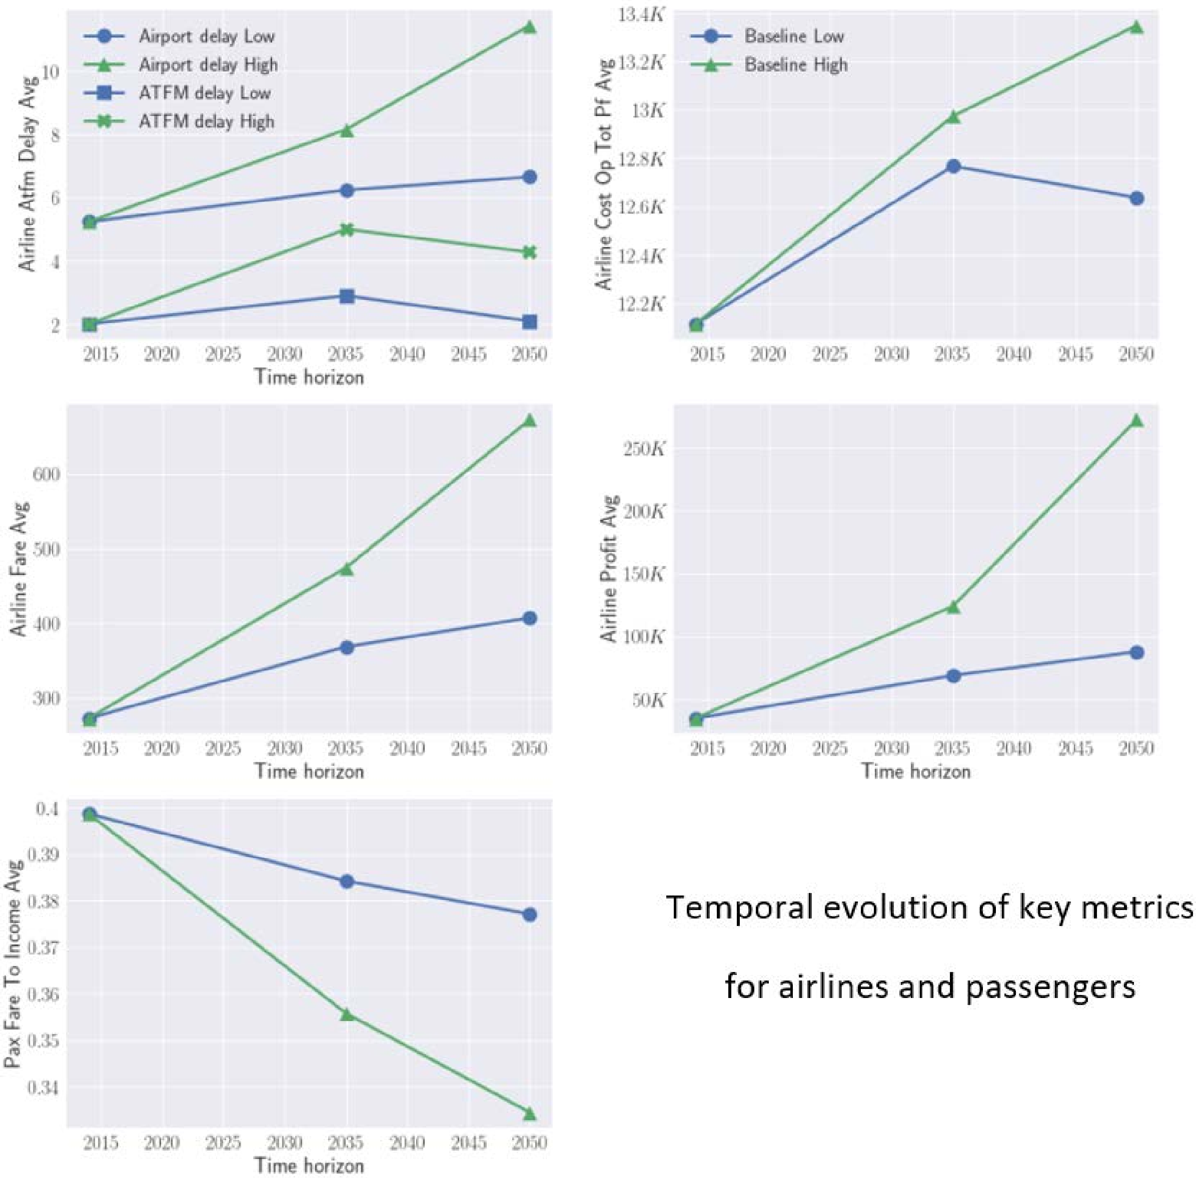

Figure 71. Temporal evolution of key metrics for airlines and passengers

5.1.3 Supportive and non-supportive scenarios

As explained in Section 4.1, the supportive and non-supportive scenarios combine several effects at the same time. In particular, from the airline perspective, three main forces shape their behaviour:

- improved technology in the supportive scenario with respect to the non-supportive;
- massive increase in the price of fuel in the non-supportive scenario;
- increased regulation schemes in the supportive scenario.

As a result, it is not trivial *a priori* to predict even the relative values of many metrics in both scenarios. To study their effect, we compare the different scenarios in 2035 first and then move on to 2050.

5.1.3.1 Supportive/non-supportive in 2035

A first look at the right-hand side of Figure 72 shows how close the scenarios are in terms of traffic. Whereas the supportive one is close to the baseline, the non-supportive is much closer to the current scenario. As a consequence, the delays display a similar pattern. Indeed, capacity is almost the same in the three scenarios, since it is mainly fixed by background factors, common to the three of them.

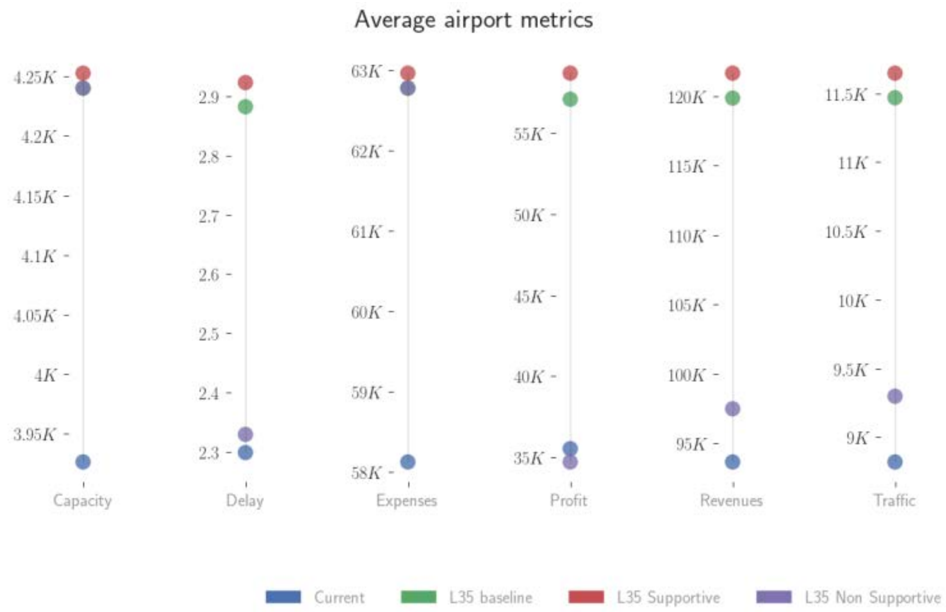

Figure 72. Average metrics for airports in different scenarios

In Figure 73, we see that the situation differs for ANSPs. Indeed, the average delay for them is only significantly different from the target delay in the baseline scenario. This is because in the supportive scenario, the traffic is very high, but so is the cap on the capacity, thanks to strong technological advancements. As a result, the traffic is manageable, and the delay is low.

The non-supportive scenario has a smaller cap, but its traffic is also much lower (actually close to the current one). The consequence is that the delay is also low in this scenario. In this scenario, the ANSPs have actually more slack capacity than in all the others. The unit rates are very close to each other in the three 2035 scenarios, which shows that the combined foreground factors do not have a strong effect on this metric, or cancel each other out.

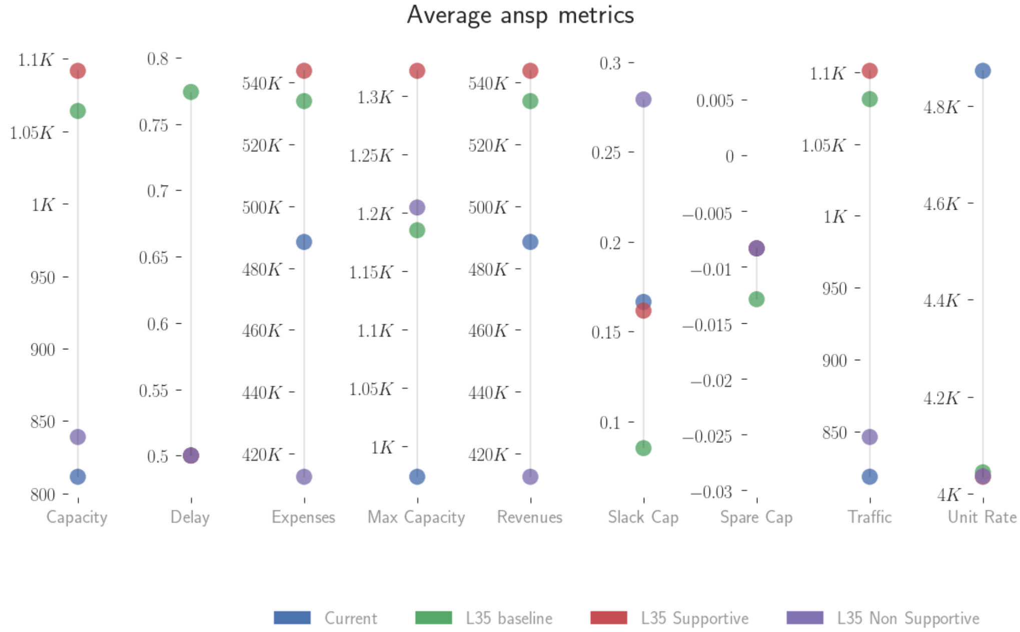

Figure 73. Average metrics for ANSPs in different scenarios

In Figure 74, we see the confirmation of the above analyses for airports and ANSPs. Moreover, we observe that the total operational cost is very high in the non-supportive case with respect to the others (twice as much), which is the reason behind the overall low traffic in this scenario. This cost is mainly driven by the augmentation in the price of fuel, leading to an increase by more than a factor five for the cost of fuel. This effect is better illustrated with an analysis of the cost structure, as shown in Figure 75. Before moving to that, note that the uncertainty has an interesting pattern too, with the non-supportive scenario faring better than the baseline in this case.

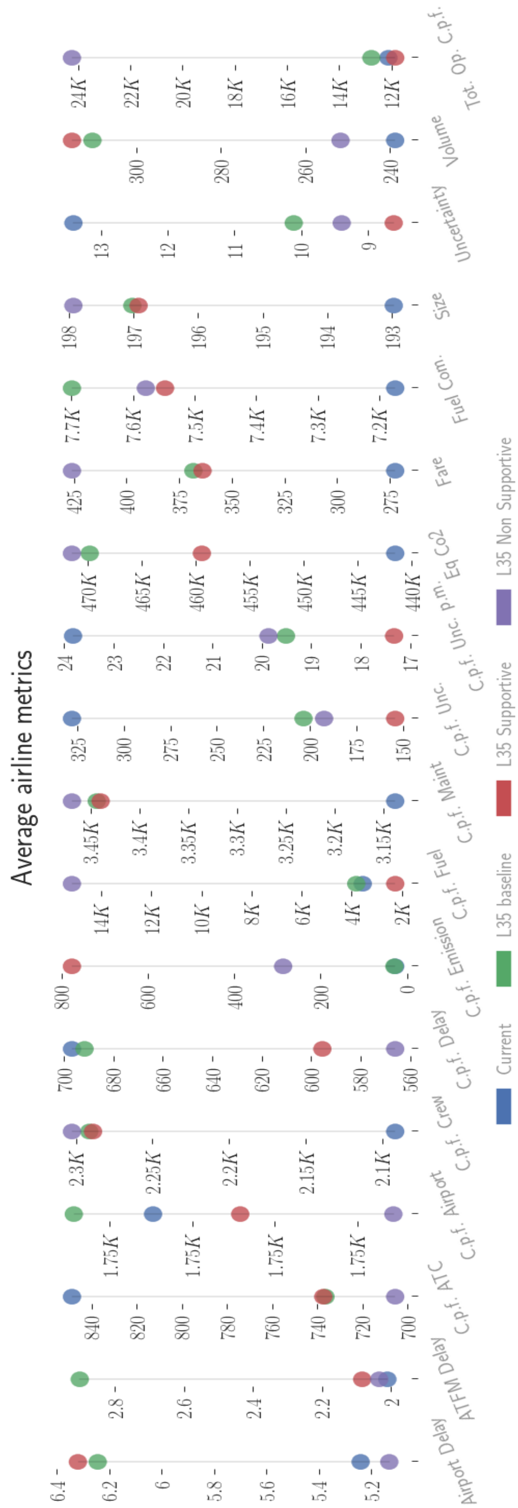

Figure 74. Average metrics for airlines in different scenarios

(Reminder: eq. CO_2 is in equivalent CO_2 of long term climate impact, ART100)

In Figure 75 we show in more detail the cost structure of the airlines. Since the cost of fuel is the most important part of the cost structure, reducing or increasing it has an immediate and sizable impact on the total cost. It is interesting to see also that the emissions cost is significantly higher in the supportive scenario, whereas it is almost absent in the baseline. This is both because NO_x emissions are important and get charged in the *supportive* scenarios (see Figure 80 in Section 5.1.4.1), and because the price of the allowances increase substantially. The variations in the cost of delay and the ATC costs are very small overall, compared to the others.

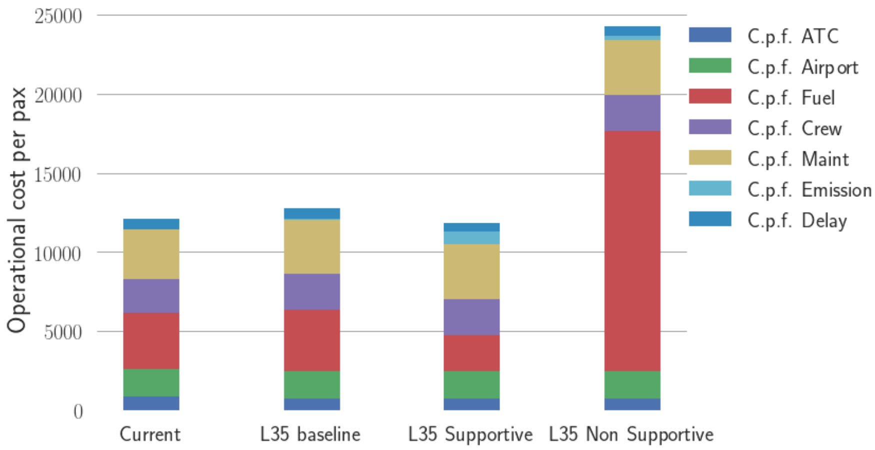

Figure 75. Cost structure of the airlines in different scenarios

Apart from the delays and volume effects already discussed above, it is interesting to see that the fares are very high in the non-supportive scenario, to the point that the fare-to-income ratio is actually higher in this scenario than in the initial situation. This is a direct consequence of the slow technological advancements and very high operational costs of the airlines.

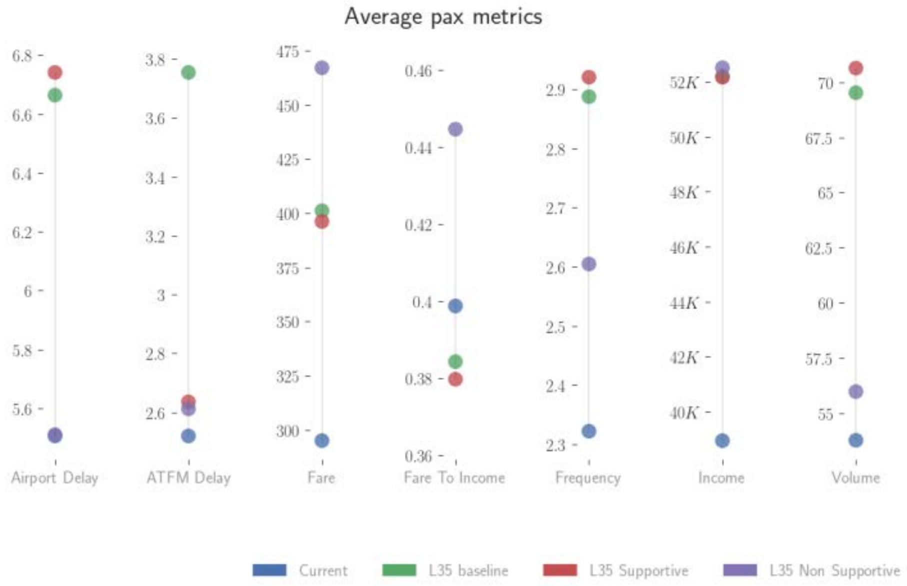

Figure 76. Average metrics for passengers in different scenarios

5.1.3.2 Supportive/non-supportive in 2050

We now move to the description of the 2050 scenarios. As for the baselines, we only highlight a few important metrics with temporal plots here. We show in Figure 77 how ANSPs create delays in the different scenarios. As shown in the top left panel, only in the baseline are the ANSPs too close to their maximum capacity (and only in 2035). In the supportive scenario, ANSPs have a much higher maximum capacity, whereas in the non-supportive case the traffic is far too low to be a problem. The situation is even better in 2050, as the gap between required and maximum capacity is very large in every scenario.

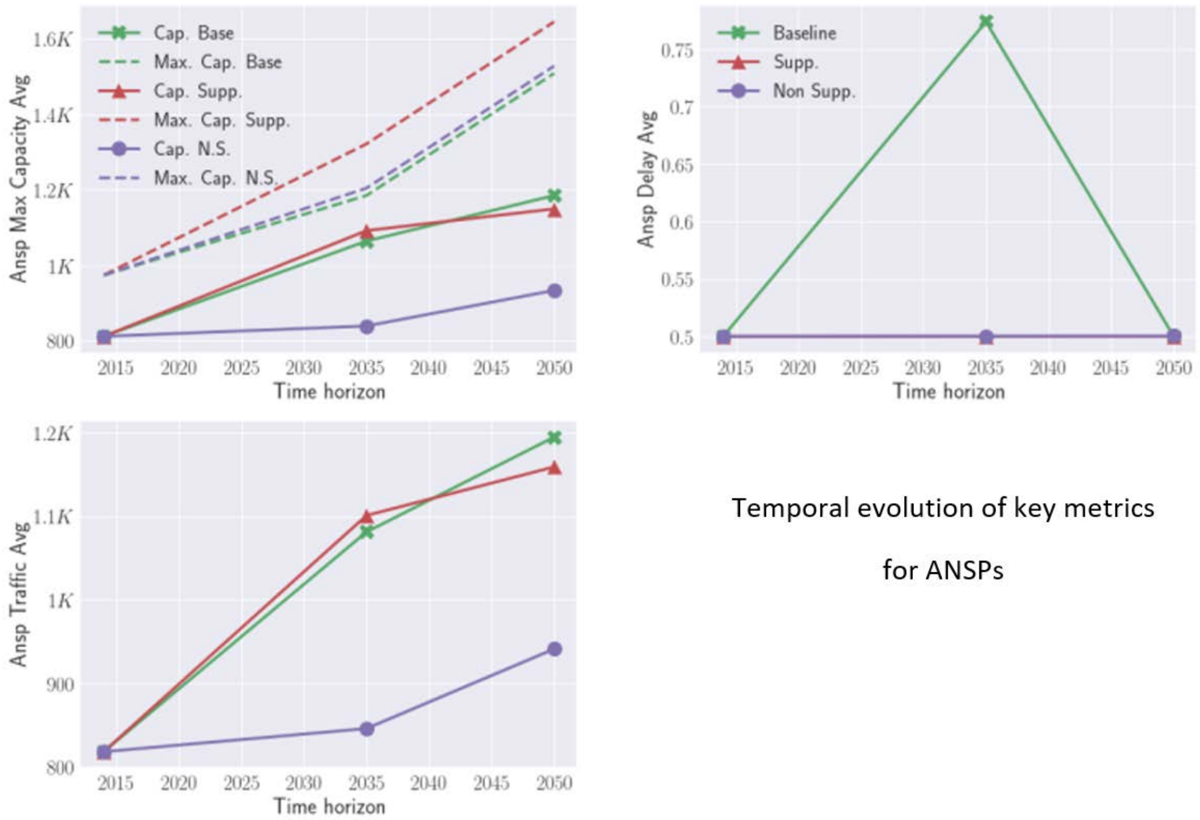

Figure 77. Temporal evolution of a few key metrics for ANSPs

In Figure 78, we show that the delays due to airports increase monotonically, contrary to the ATFM delays. Moreover, the supportive scenario is very close to the baseline, whereas the non-supportive case is rather less, because of lower levels of traffic. This does not stop the operational cost from being vastly superior in the latter, but interestingly, the supportive scenario is also above the baseline in 2050, contrary to 2035. This is mainly due to increasing emissions cost, as shown on the bottom-left graph. Higher allowance costs and heavier aircraft combine to produce a significant cost, now comparable to the other costs. The fare-to-income ratio decreases in 2035 for the supportive scenario, but slightly goes up again in 2050, contrary to the baseline. This is probably due to the increase of the operational cost due to emissions costs.

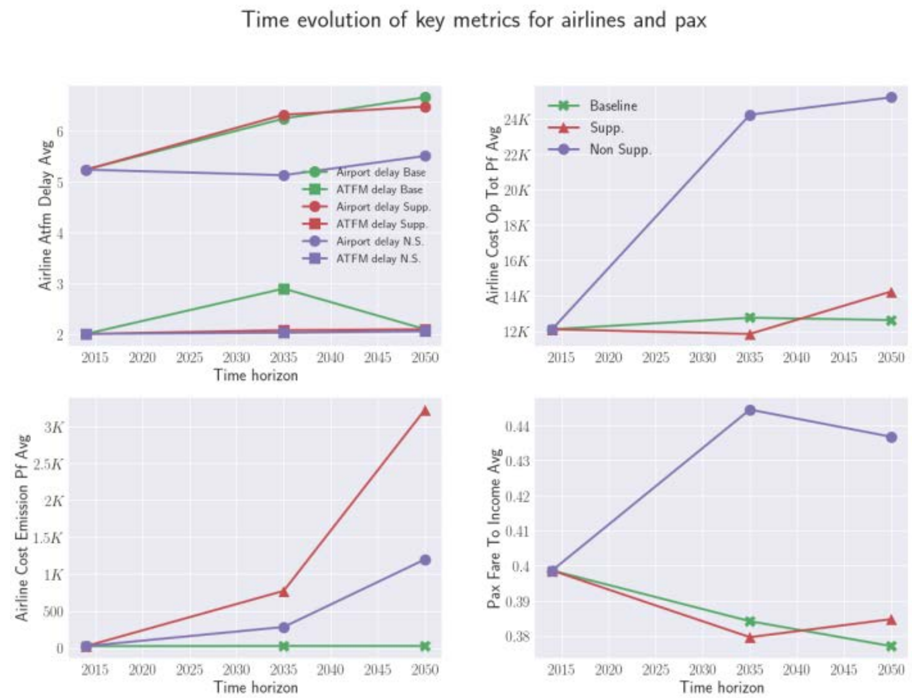

Figure 78. Temporal evolution of few key metrics for the airlines and passengers

5.1.4 Summary of key strategic results

5.1.4.1 Key metrics

In order to summarise the main results of the model, we present a set of key metrics that concern all five stakeholders (airports, airlines, ANSPs, passengers, environment). These metrics have been chosen for their relevance and/or for their originality with respect to similar studies:

- the cost of delay has been chosen for its importance in airline operations; it is very rare to have an estimation of this cost taking the full distribution of delays into account;
- the cost of uncertainty, part of the above cost, is even more rarely measured; the fact that SESAR envisions substantial reductions of uncertainty in the future, calls for a proper estimation of the associated benefits;
- the total level of emissions in CO_2 equivalents has an obvious importance, given rapid climate change; the emissions are all too often measured per flight, which does not show their full impact; note that not only does the model predict a massive increase in the total emissions, but also an increase in emissions per flight, since heavier flights are expected to fly in the future;
- the fare-to-income ratio is important to properly estimate whether the system is more efficient economically or not; as noted above, this metric is probably underestimated in the model; however, it is interesting to see that it still decreases in some cases, even without strong competition between airlines;

- the total delay per flight shows how much the system will be under stress in the future; if the airports are indeed the main contributors to the delay, we show that the ANSPs have a high potential for greater delays, since they move close to their maximum capacities, even with the technological advancements envisioned by SESAR.

In Figure 79, we show the evolution of these key metrics. The total cost of delay includes the compound effects of the mean delay and the uncertainty (shown second, alone). The total emissions in CO_2 equivalents is for the entire system, taking into account CO_2, NO_x, H_2O and contrails (the last two not being charged in any scenarios). The last plot is the sum of the delay produced by ANSPs and the delay produced by airports, per flight.

The cost of delay and the cost of uncertainty decrease substantially in both the supportive and non-supportive scenarios, significantly more than the baseline for the cost of delay. On the other hand, we can see that the total CO_2 equivalent emissions increases substantially, and that total delay increases, especially in the supportive scenario. The fare to income ratio decreases slightly in the supportive case, as well as in the baseline.

Figure 79. Key metrics evolution

Note that the emissions in the figure are the total emissions, summing the CO_2, NO_x, H_2O, and contrail effects, but it is interesting also to see which type of emissions has the bigger impact. We show in Figure 80 that the biggest greenhouse effects are not due to CO_2, but rather to NO_x, which is not yet charged by any scheme. According to the model provided by DLR, contrails also have a very substantial impact. Water has a lower impact. Note that this partition does not change much across scenarios.

This result, together with the high increase in the total emissions, means that charging for the CO_2 only is not enough and that it is urgent to consider all types of emissions at the same time when devising an environmental scheme.

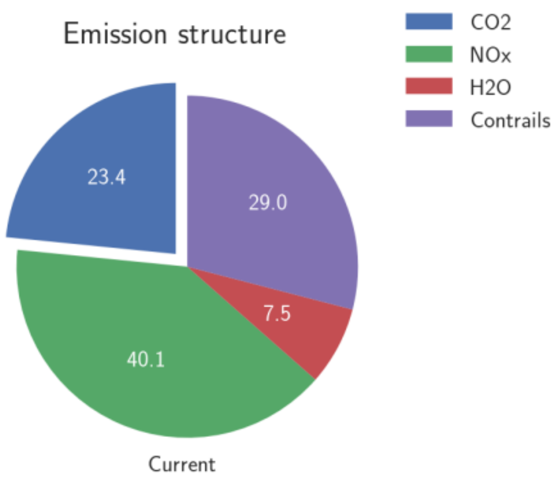

Figure 80. Emission structure

(Reminder: eq. CO_2 is in equivalent CO_2 of long term climate impact, ART100)

5.1.4.2 Key results

We conclude this analysis by highlighting the main points that emerge from the strategic results.

1. The main drivers for most of the metrics are:
   a. the demand (set in the model by the average income of passengers);
   b. the price of fuel (as shown also more in detail in D5.1).
2. ANSPs may get close to their maximum capacity (set to 120% of current capacity, with technological advancements), depending on the scenarios, and trigger some significant delay for airlines.
3. ANSPs see their unit rates decrease substantially in the future. This is due to the joint effect of higher levels of traffic and greatly improved efficiencies.
4. Airports create most of the delay, and the increase of capacity envisioned by SESAR is not sufficient on its own to deal with the increase in traffic.
5. The cost of emissions only really has an impact on airlines when **NO_x is taken into account** together with a great increase in the price of allowances. Emissions have otherwise (almost) no impact on the cost structure of the airlines.
6. The average size of the aircraft used by airlines is increasing. This has an impact on the average cost of fuel and environmental impact per flight, despite various other improvements such as the length of flights. Other technological improvements related, for instance, to engine efficiency have not been considered in Vista.
7. Total emissions are expected to soar in the future. This is mainly driven by the increase in traffic, and to a lesser extent by the increase in the average size of the aircraft, see previous point.

8. The reduction of uncertainty envisioned by SESAR is expected to a have major impact on the cost of delay to the airlines. The cost of uncertainty represents roughly half of the total cost of delay in the current scenario, but its share drops substantially in the future.
9. Passengers usually see a moderate decrease of the fare with respect to their income, except when the operational cost of the airlines increases too much, notably because of the price of fuel. This drop could be larger if the competition were to be stronger in the model, since airlines seem to make too much profit.

Among these qualitative main results, #6 is important to note as there are some conflicting studies about the increase of the size of aircraft in the future. It is interesting to see that we obtained this result without actually changing the size of the aircraft anywhere, but just by allowing the most profitable flights to grow faster than the others, which leads naturally to a higher average for the size.

The impact of uncertainty (#8) on the ATM system is also an open question. In the model, we have only used unpredictability to compute real expectations of costs as opposed to costs of average delay, as is often the case. With this simple setup, we obtain the important result that the uncertainty counts for around 50% of the cost of delay in the current scenario. As a result, any changes to uncertainty substantially impacts the cost of delay and thus the total operational cost of the airline (note however that the cost of delay is not the major contributor to the operational cost of the airline). The changes envisioned by SESAR – although optimistic – thus will have a great impact on the operations. Note also that we did not take into account other phenomena linked to uncertainty: change of buffers, suboptimal human decisions due to time constraints, etc.

The model also answers an important question regarding the ANSPs (#2, #3). Keeping the current scheme, and provided that the SESAR innovations are implemented, we find that the system is able to cope with the increase of traffic in most scenarios, at least in the long run, but is dangerously close to its limit. Moreover, the gains in efficiency and the gains in capacity mean that the ANSPs are able to reduce their unit rates very significantly, driving the ATC costs down for the airlines. The validity of this result is also subject to whether the maximum capacity of the ANSPs is actually currently around 120% of their capacity on average. Moreover, the fact that ANSPs are currently differently close to their limits could have an impact on the results.

Airports are not able to cope with the increase of traffic with the SESAR improvements alone. They already produce a large share of the delay currently, and are likely to produce most of it in the future. Several other factors mentioned in D3.1 could help them (e.g. better passenger processing) but have not been implemented. Infrastructure expansions seem, however, to be the best option for many of them.

The main trade-off highlighted by this study is the environment *versus* the increase in traffic. It has been known for years that economic development is prejudicial towards climate impact, in the short and medium term. We find this to be the case in aviation also. Even taking into account the fact that we did not include increased efficiency of aircraft engines, it is crucial to note that the overall efficiency of the system is nowhere near enough to offset the potential increase in traffic, let alone decrease the emissions. Even more important, we highlight the fact that gains in efficiency in the model turn into increased profits, and thus into the expansion of operations, which negate the gains *in fine* from an environmental point of view.

5.2 Pre-tactical layer

The pre-tactical layer is translates the flight schedules into flight plans, the passenger flows into passenger itineraries, and the changes in demand and capacity into changes in probability of incurring delay due to ATFM regulations. Therefore, the results presented for the pre-tactical layer focus on metrics related to the flight plans, passenger itineraries and ATFM delay.

The metrics that are quantified in the pre-tactical layer include:

- For flight plans:
  - distance (NM);
  - (flight) time (minutes);
  - en-route airspace costs (euros);
  - estimated fuel and carbon costs (euros);
  - estimated flight plan fuel usage (kg);
  - buffer and taxi times (i.e., block time - flight plan time) (minutes).
- For NASs:
  - revenue due to airspace en-route charges (euros);
  - demand (aircraft entry counts);
  - probability of assigning delay due to ATFM capacity regulations;
  - estimated total amount of delay generated.
- For passengers:
  - total number of passengers;
  - load factors;
  - number of connecting passengers.

In Deliverable D5.1, the usage of the pre-tactical layer to analyse the changes in flight plan selection and its impact on the demand and revenues per ANSP, cost per airlines, etc., where presented for scenarios where the fuel cost was modified with the 2014 demand and for scenarios where the airspace en-route charges were modified. These results can be found in D5.1, the results presented here will focus on the scenarios described in Section 4.

5.2.1 Baseline evolution

As with the strategic layer, the first analysis is performed with respect to the baseline defined by the current, low 2035 baseline, and the low 2050 baseline scenarios.

5.2.1.1 Low baseline

(i) Flight plans

Figure 81 presents the average metrics to compare the impact of the scenario on the flight plan characteristics. The metrics that are presented are the flight plan distance, flight plan time, buffer time (computed as the difference between SOBT and SIBT and the flight plan time, i.e., the time allocated strategically for taxi in and out and arrival buffer), the fuel of the flight plan and the associated costs (fuel, CO_2 and airspace en-route charges). For each flight there are usually different flight plan options to choose from, each of those options will have different characteristics. We can, therefore, analyse on average how much longer the selected flight plans are with respect to the shortest available, and how much more expensive they are for the en-route charges and for the fuel costs, among other possible comparisons. Note that it might not be possible to optimise all the metrics at the same time as the flight plan that minimises one of them does not necessarily minimises the other parameters, thus leading to trade-offs. An example of this is when longer routes use cheaper airspaces.

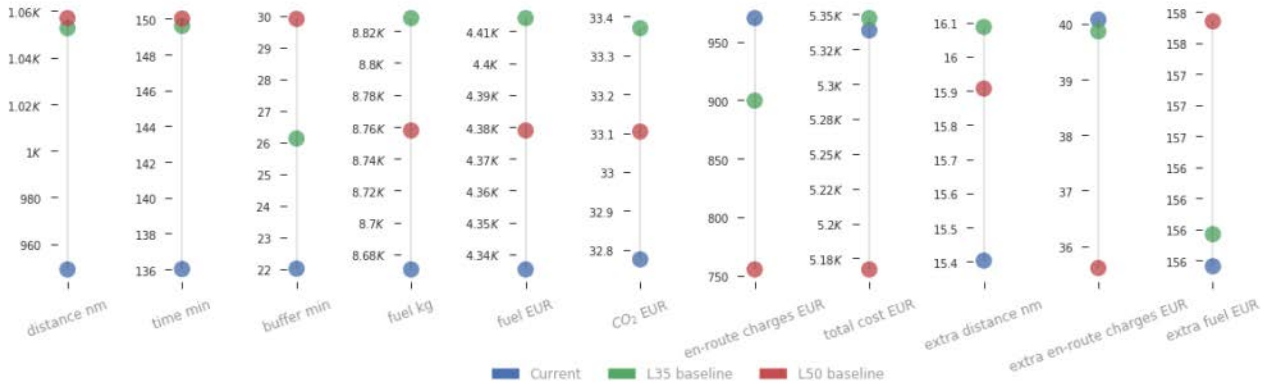

Figure 81. Average metrics for flight plans in low baseline

For the low baseline, we can observe that for 2035 and 2050 the flight plan distances are on average longer than in the current scenario. The flying times are also longer. However, there does not seem to be a great difference between 2035 and 2050 for these two indicators. Interestingly, the time allocated for taxi in, taxi out and flight plan buffer increase on average, progressively from current to 2050, starting at 22 minutes and increasing to 30 minutes. A possibility to explain the fact that buffer times increase between 2035 and 2050, but the flying time and the flight plan distance remain at similar values, might be that some of the factors are reducing the flight plan distance (affecting the ENV SESAR KPI) and, even if flights have on average the same flight plan distance and flight plan time, in reality, airlines are operating longer origin and destination pairs that have larger buffers in their schedules. Note also how the fuel consumption is on average higher for 2035 than for 2050, translating directly to higher fuel costs. Airspace en-route charges tend, however, to decrease per flight.

Comparing the shortest and cheapest routes possible, it is interesting to note that in 2035 and 2050 the routes are slightly longer than for the current scenario, but the en-route charges are closest to the cheapest option available even though that means selecting flight plans that on average have a slightly higher fuel consumption than the most efficient flight plan from a fuel consumption point of view, compared with that in the current scenario.

Table 41 presents the average MTOW of the aircraft in each scenario and the average number of flight plan options per schedule. Note how there is a trend to have larger aircraft to accommodate the growing demand.

Table 41. Average MTOW and number of flight plan options low baseline

Scenario	MTOW	Flight plan options
Current	85.7	8.9
2035 baseline low	86.5	14.0
2050 baseline low	87.5	9.8

(ii) ANSPs

We can compute the traffic demand per ANSP for each scenario to see how the scenario affects the amount of traffic and its temporal evolution. An example of this was presented in Figure 13.

Vista is also able to compute the total revenues due to en-route airspace charges per ANSP, as shown in Table 42, and the variation per ANSP, as presented in Figure 82. Note that in these radar plots, the values are normalised per axis, i.e. 1 is the highest revenue obtained for a given ANSP.

Table 42. Total en-route airspace revenue M EUR low baseline

Scenario	En-route airspace charges revenue M EUR
Current	26.2
2035 baseline low	30.6
2050 baseline low	28.1

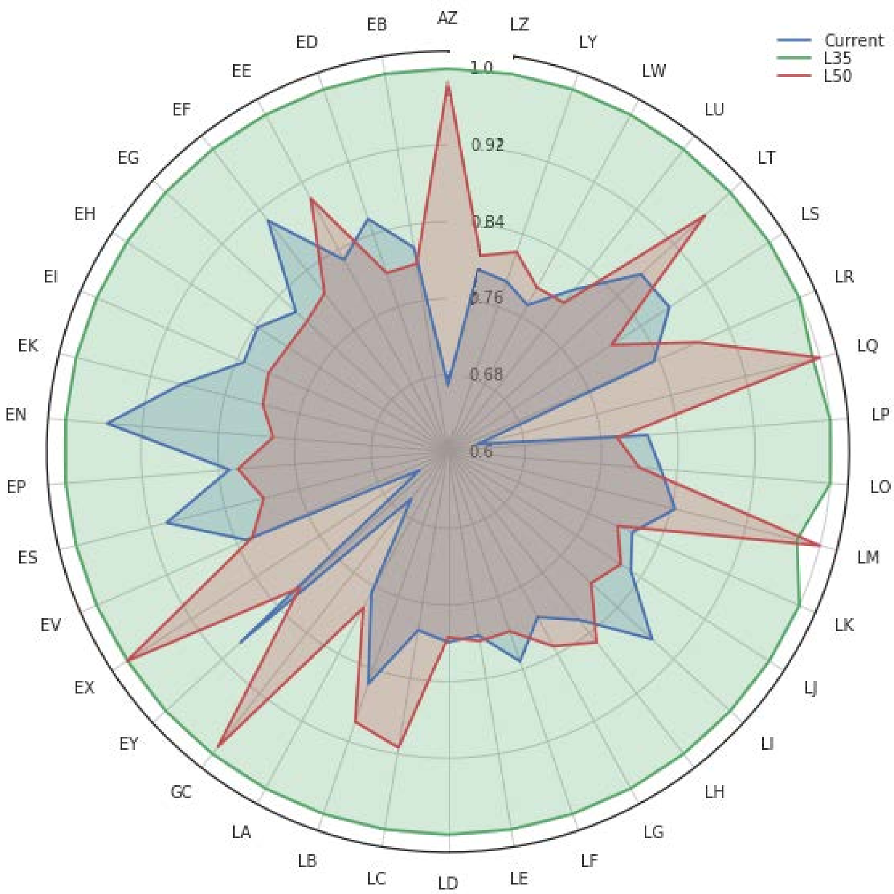

Figure 82. Revenues due to en-route charges variation per ANSP low baseline

Note that in the case of the low baseline comparison, the revenues between current and 2050 do not change significantly even though some ANSPs might have large variabilities (e.g., EN reduces its revenues by more than 10%, while the airspace of Azores (AZ) is used more). The 2035 baseline case is the one that for all the ANSPs, except LM, represents the highest revenues.

We can compare the probability of having ATFM delay due to capacity regulations. Figure 83 shows how the probability of having ATFM delay due to capacity regulations tends to reduce from current to 2035 and 2050 even if not substantially. This is due to the fact that the improvements in the system are offset by the increment in demand.

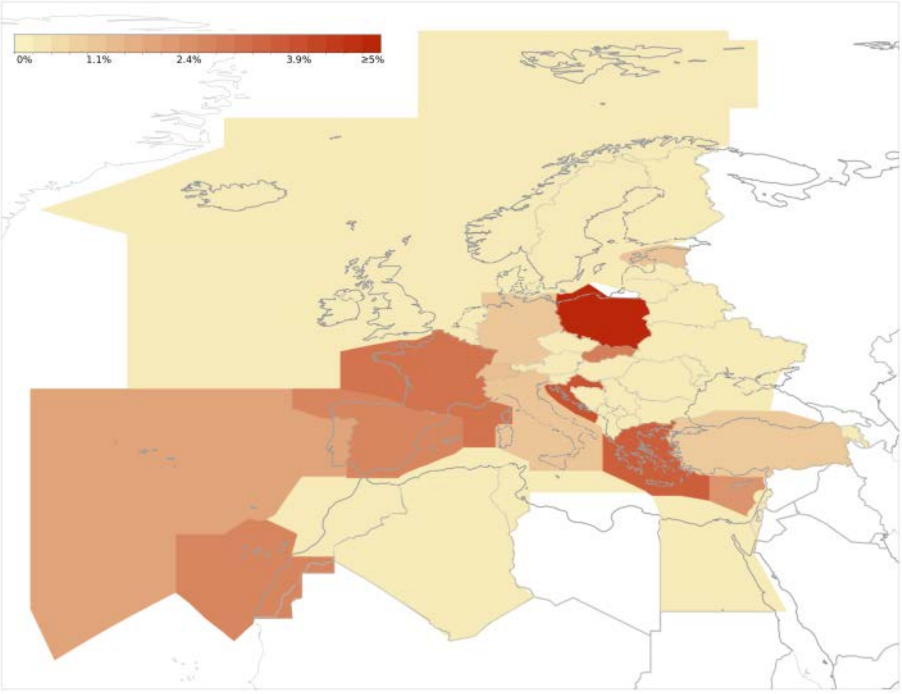

a) Current

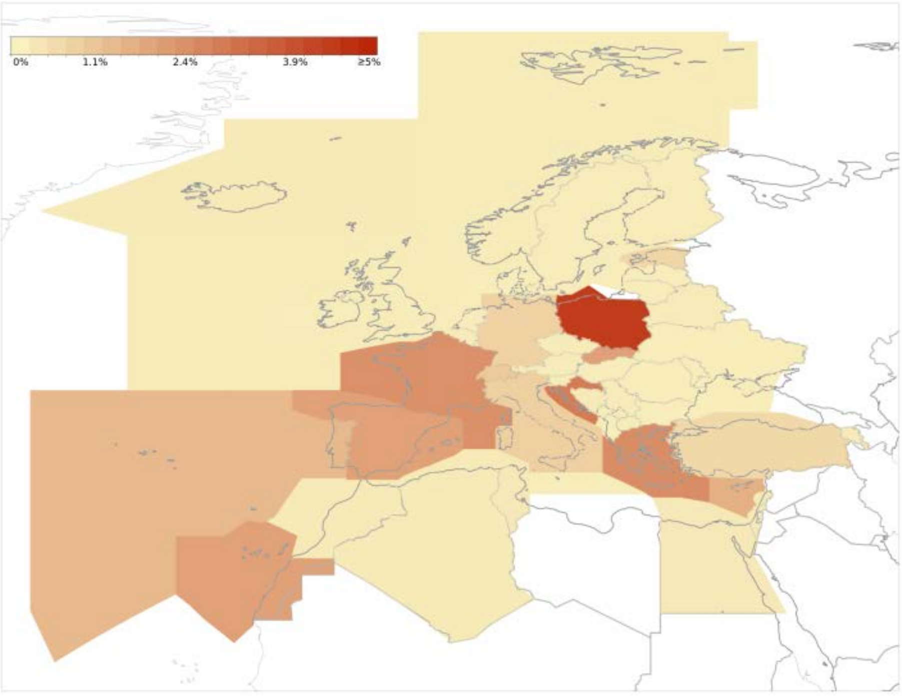

b) 2035 baseline low

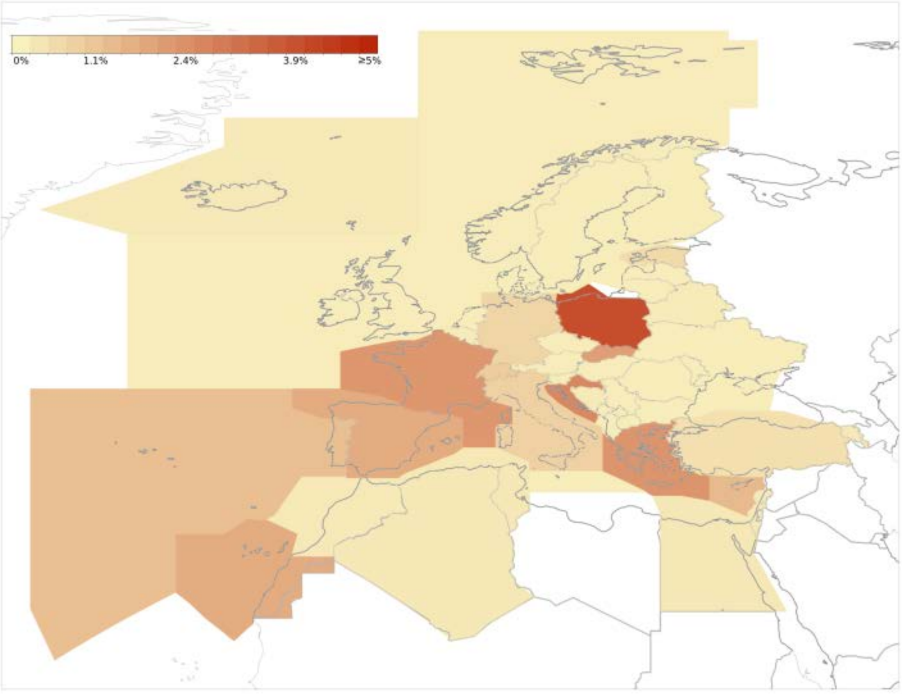

c) 2050 baseline low

Figure 83. Probability ATFM delay capacity low baseline

(iii) Passengers

Table 43 shows the total number of passengers and the number of connecting passengers for the scenarios in the low baseline cases. The number of passengers between 2035 and 2050 does not increase substantially.

Table 43. Passenger details (per day), baseline low

Scenario	Total number passengers ('000)	Passengers with 1 connection ('000)	Passengers with 2 connections ('000)	Total passengers with connections ('000)
Current	3 401	275 (8.1%)	9 (0.3%)	284 (8.4%)
2035 baseline low	4 344	373 (8.6%)	12 (0.3%)	385 (8.9%)
2050 baseline low	4 740	412 (8.7%)	13 (0.3%)	425 (9.0%)

5.2.1.2 High baseline

The analysis of the high baseline for 2035 and 2050 allows us to see the impact of a high economic and technological development and how this would affect the flight plan, ANSPs, and passengers' itineraries.

(i) Flight plans

As shown in Figure 84, the distance flown, flight plan time and fuel usage tend to increase over time, as in the low baseline case. The buffer times also increase and, in this case, are larger than in the low case. This could be due to the fact that the background factors are reducing the flight plan distance needed to cover larger distances between origin and destination pairs. Also, as observed in Table 44, the average MTOW of the aircraft used increases significantly. This could be related to using larger aircraft as demand increases. An effect of this increment in aircraft size might represent more flights covering longer routes, which have larger buffers in their schedules.

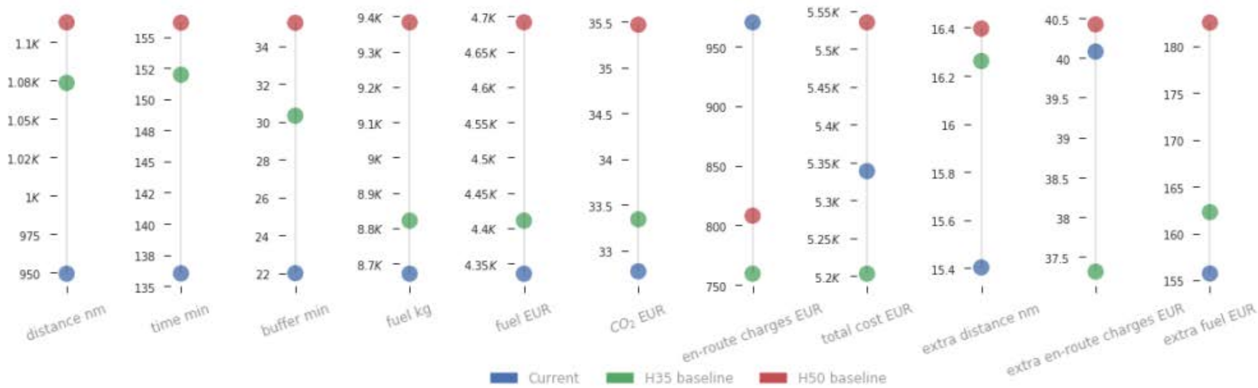

Figure 84. Average metrics for flight plans in high baseline

The en-route airspace charges on average per flight tend to decrease, particularly for the 2035 high baseline scenario. This is linked to a decrease in unit rates provided as an outcome of the economic layer and due to the increment in demand, which spreads the increment in cost among more users. For the 2050 high scenario, the rates do not fall as much as for the 2035 case, due to the increment in costs to provide more capacity.

Table 44. Average MTOW and number of flight plan options high baseline

Scenario	MTOW	Flight plan options
Current	85.7	8.9
2035 baseline high	89.9	14.8
2050 baseline high	98.5	10.5

Table 44 also presents an increment in the average number of flight plan options per flight schedule, which is consistent with a larger use of longer flights, which have more possible routes available.

(ii) ANSPs

For the ANSPs the high baseline scenarios represent a large increment in the revenues obtained. This is particularly true for the 2050 baseline high case when revenues nearly double c.f. the current scenario. The cost per flight, however, as presented in the previous section (see Figure 84), decreases. This means that the increment in traffic allows a reduction in the unit rate, while obtaining more revenues, which offset the increment in cost to provide the extra capacity.

Table 45. Total en-route airspace revenue M EUR high baseline

Scenario	En-route airspace charges revenue M EUR
Current	26.2
2035 baseline high	33.3
2050 baseline high	45.1

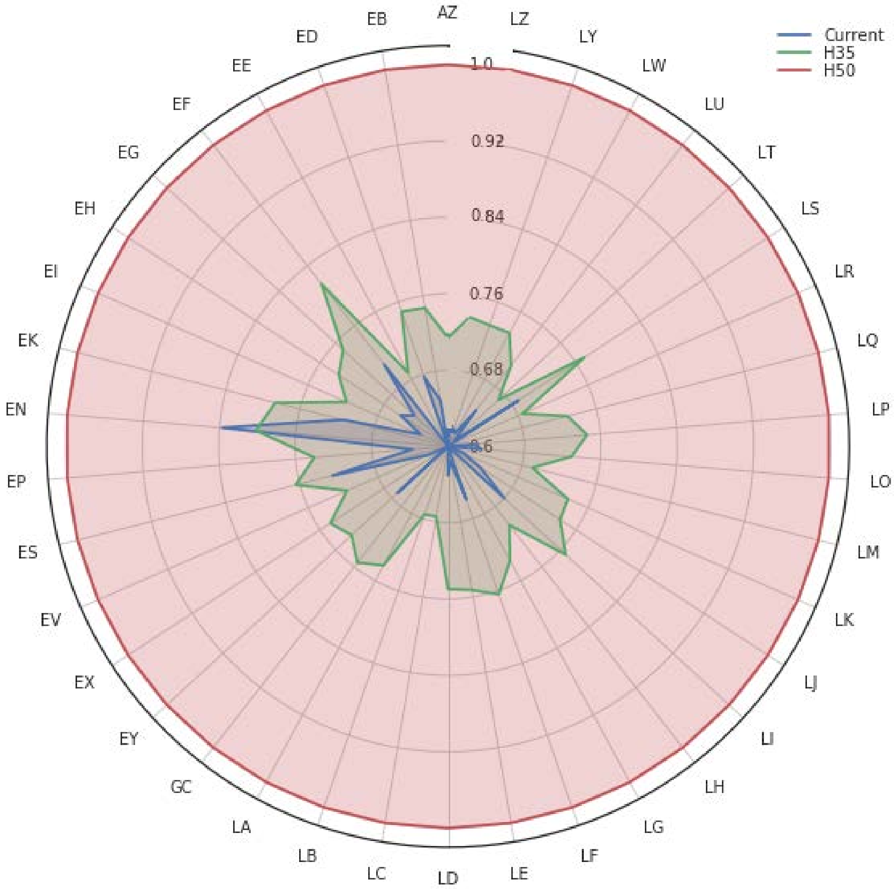

Figure 85. Revenues due to en-route charges variation per ANSP high baseline

It is clear from Figure 85 how the high 2050 scenario represents an increment of more than 30 percentage points for all the ANSPs. The 2035 or current values are around 0.65 - 0.75, i.e. approximately 30-40% lower than the 2050 case, as per the normalisation described above.

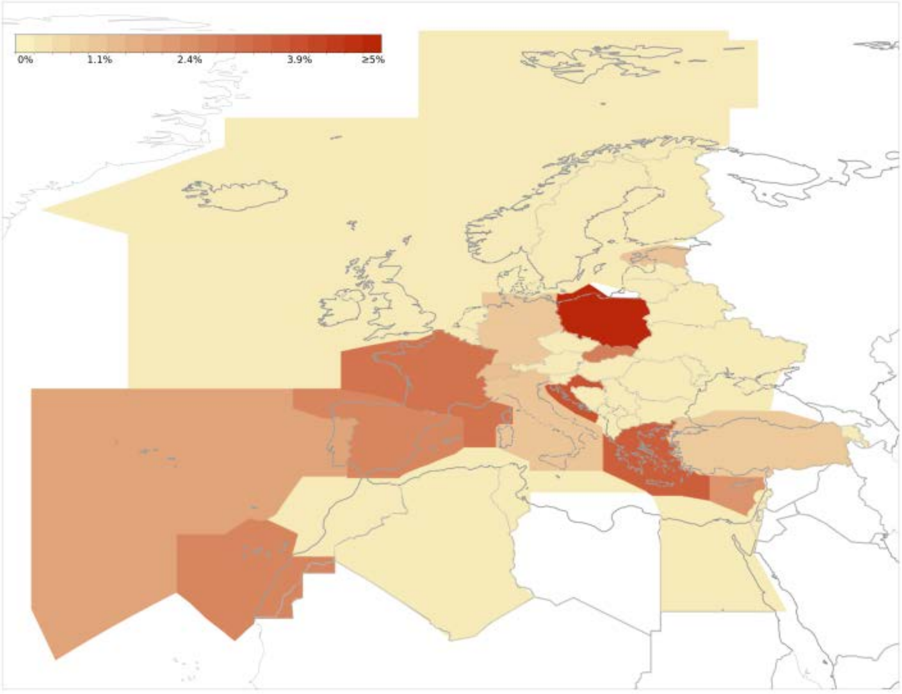

a) Current

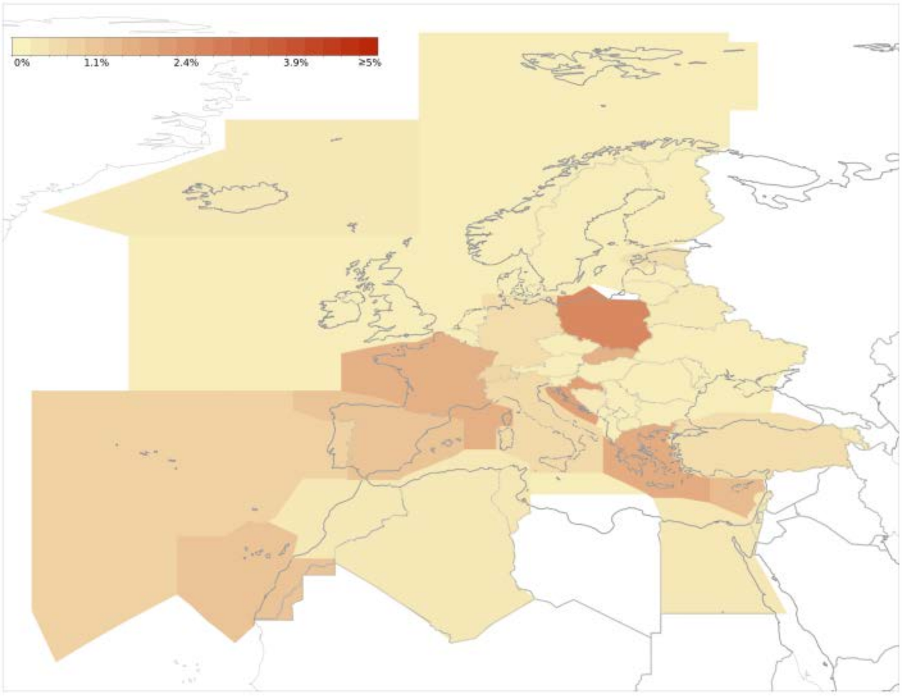

b) 2035 baseline high

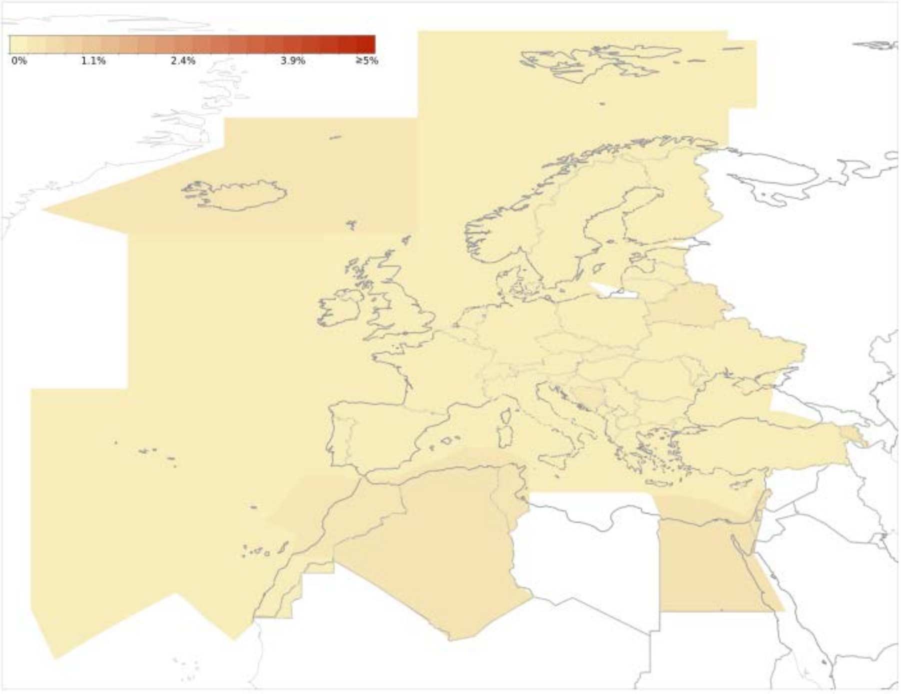

c) 2050 baseline high

Figure 86. Probability ATFM delay capacity high baseline

Looking at the probability of incurring delay due to ATFM capacity regulations, see Figure 86, we can observe that these probabilities decrease over time. This could be counter intuitive as the traffic is growing. However, demand is increased thanks to the technological developments which counterbalance the extra pressure on demand in the system, arriving at a situation whereby the flights delayed due to capacity restrictions are lower over time.

(iii) Passengers

Regarding the passengers, the first observation from the following table is that there are significantly more passengers in 2035 and in 2050 than in the current scenario. This increment is noticeable also with respect to the baseline low scenarios (presented in Table 43), for example in the 2050 case there are over 7.7 M passengers in the high scenario and 4.7 M in the 2050 low baseline case. As observed in the low scenarios, the increment in number of passengers is also correlated with an increment in the percentage of connecting passengers.

Table 46. Passenger details (per day), baseline high

Scenario	Total number passengers ('000)	Passengers with 1 connection ('000)	Passengers with 2 connections ('000)	Total passengers with connections ('000)
Current	3 401	275 (8.1%)	9 (0.3%)	284 (8.4%)
2035 baseline high	5 687	500 (8.8%)	15 (0.3%)	515 (9.1%)
2050 baseline high	7 753	722 (9.3%)	21 (0.3%)	743 (9.6%)

5.2.2 Supportive and non-supportive scenarios

As with the strategic layer, the analysis of the supportive and non-supportive scenarios is done first for the 2035 cases, and then for 2050.

5.2.2.1 Supportive/non-supportive in 2035

(i) Flight plans

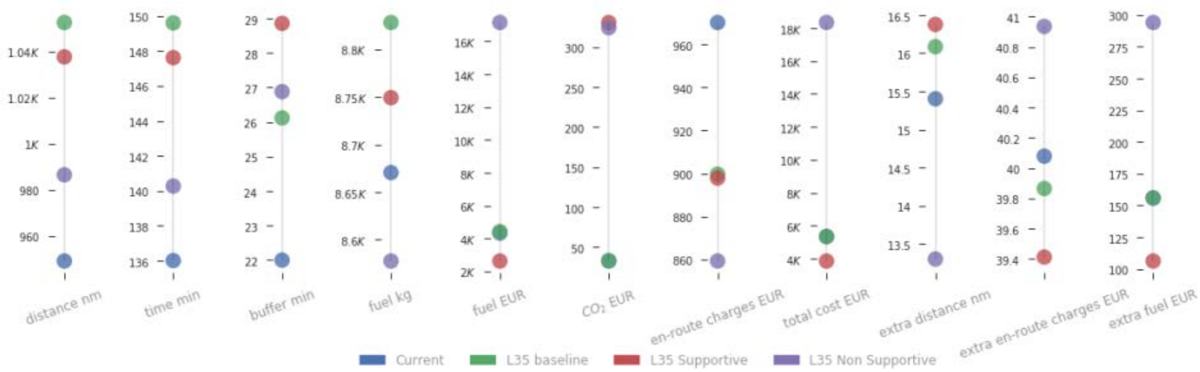

Figure 87. Average metrics for flight plans in low baseline

The results for the flight plans are described in Figure 87. There are several characteristics of the results that are worth noting. The average amount of fuel used per flight decreases with respect to the 2035 baseline, for both the supportive and the non-supportive scenarios, but the change is not very large. The distances flown are also lower for both cases with respect to the baseline, which might be related to the impact of the factors of the efficiency of the route. Flight plan distance, time and buffers are larger for the supportive scenario than for the non-supportive case. The increment in buffer times for all the 2035 scenarios with respect to the current is consistent with the usage of longer routes and larger aircraft (see Table 47).

Table 47. Average MTOW and number of flight plan options 2035 supportive/non-supportive

Scenario	MTOW	Flight plan options
Current	85.7	8.9
2035 supportive	86.8	14.2
2035 non-supportive	86.3	12.9

(ii) ANSPs

For the ANSPs, the 2035 supportive scenario is the one yielding the higher revenue due to en-route airspace charges (see Table 48).

Table 48. Total en-route airspace revenue M EUR 2035 supportive/non-supportive

Scenario	En-route airspace charges revenue M EUR
Current	26.2
2035 supportive	30.9
2035 non-supportive	25.0

As expected, the higher demand, and therefore capacity required, means that all ANSPs obtain an increment in their revenues in the 2035 supportive scenario. However, it is interesting to note that in the non-supportive case the revenues decrease with respect to the current scenario for most of the ANSPs (see Figure 88).

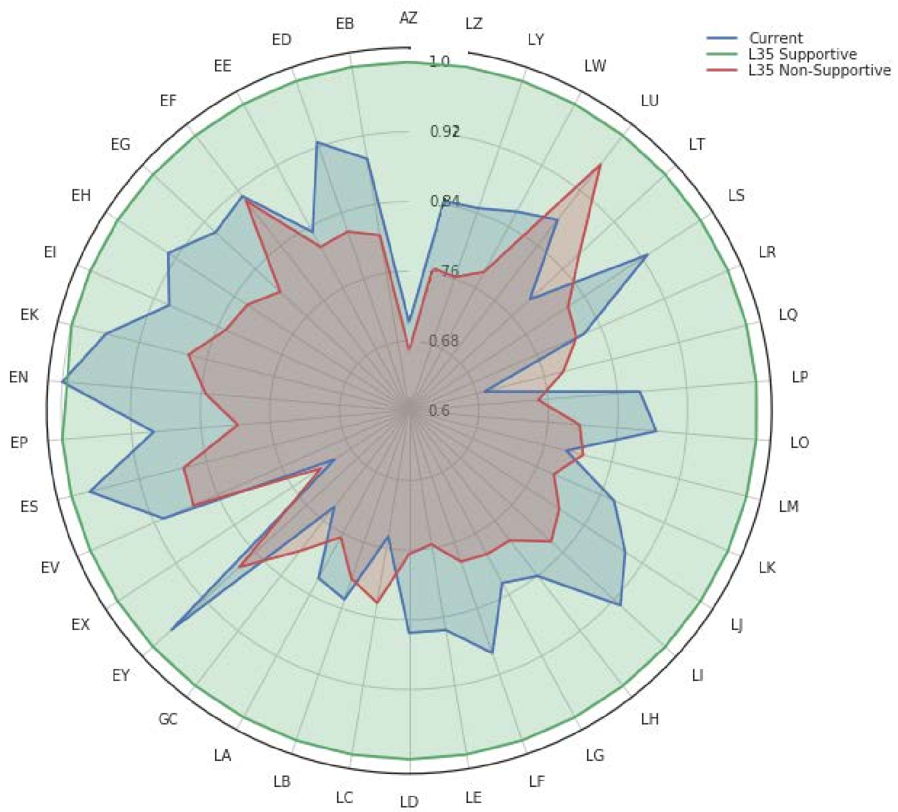

Figure 88. ANSP revenues due to en-route charges variation, 2035 supportive/non-supportive

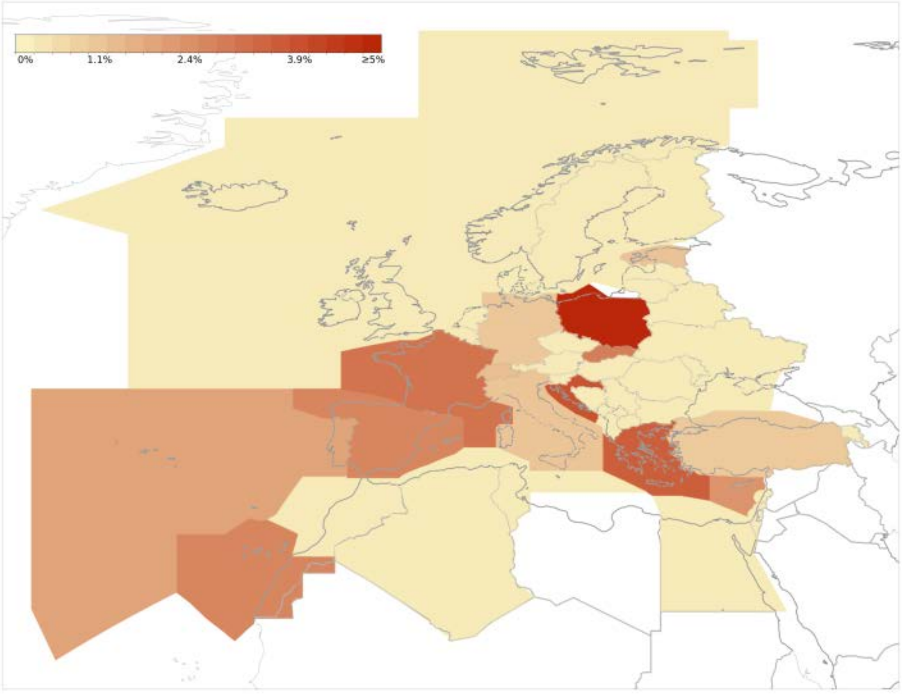

a) Current

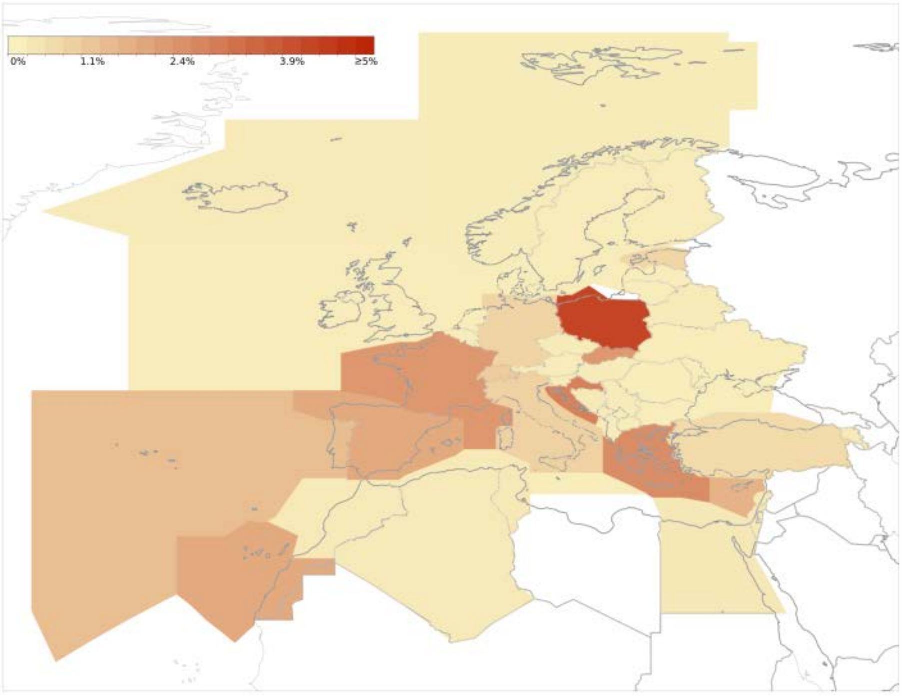

b) 2035 supportive

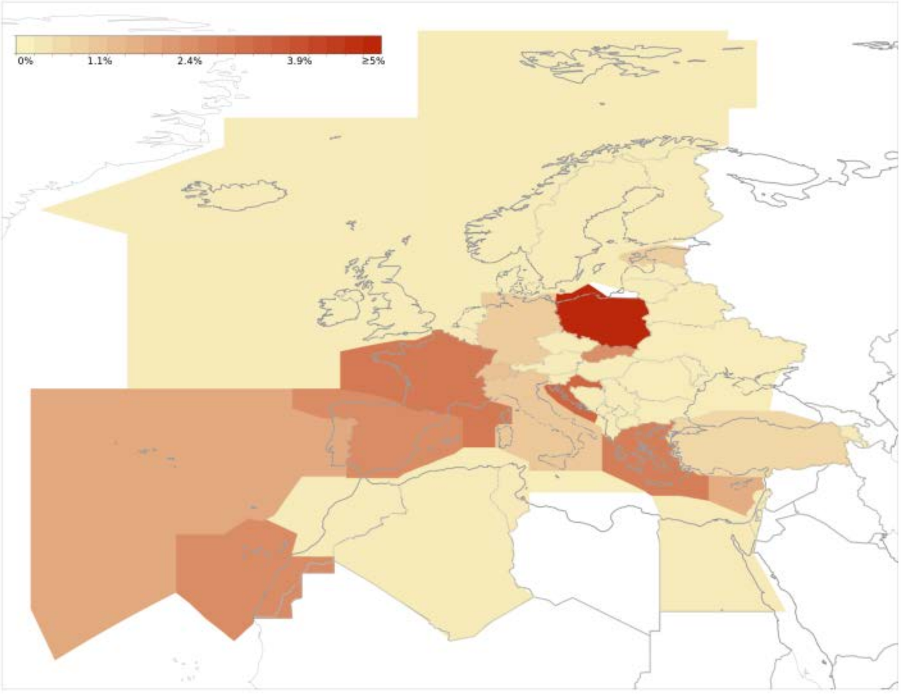

c) 2035 non-supportive

Figure 89. Probability ATFM delay capacity 2035 supportive/non-supportive

For the probabilities of having ATFM regulations due to capacity, in both cases (2035 supportive and 2035 non-supportive) there is a reduction with respect to the current scenario. However, as expected, the reduction is lower for the non-supportive case (see Figure 89).

(iii) Passengers

The increment in number of passengers is not very large with respect to the current scenario for the 2035 non-supportive case and similar in the supportive case to the 2035 baseline low, as reported in Table 43 and Table 49.

Table 49. Passenger details 2035 supportive/non-supportive

Scenario	Total number passengers ('000)	Passengers with 1 connection ('000)	Passengers with 2 connections ('000)	Total passengers with connections ('000)
Current	3 401	275 (8.1%)	9 (0.3%)	284 (8.4%)
2035 Supportive	4 395	377 (8.6%)	12 (0.3%)	389 (8.8%)
2035 Non-supportive	3 589	328 (9.1%)	11 (0.3%)	339 (9.4%)

5.2.2.2 Supportive/non-supportive in 2050

(i) Flight plans

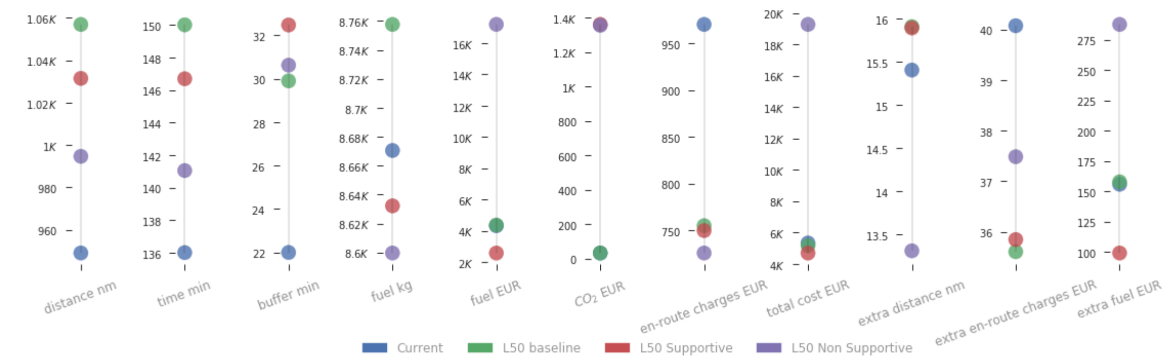

Figure 90. Average metrics for flight plans in 2050 supportive/non-supportive

The results presented in

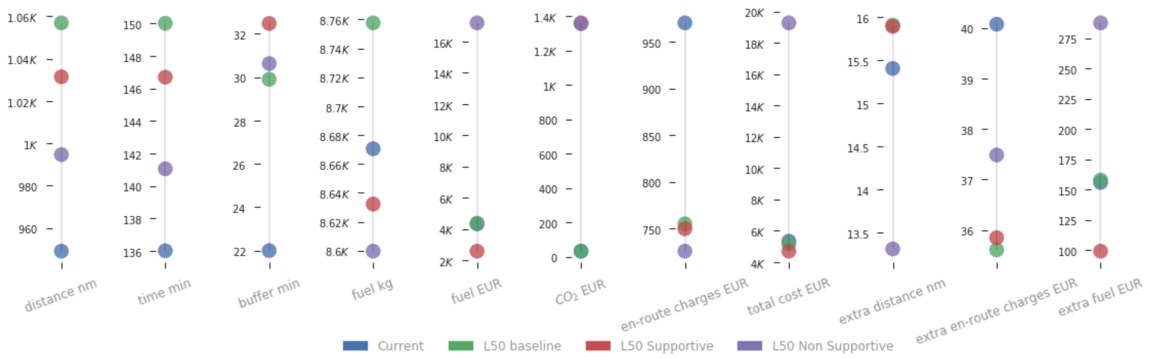

Figure ***90*** show that across the 2035 supportive and non-supportive scenarios, and their equivalents in 2050, several metrics remain very similar (e.g., flight plan distance, flight plan time, and average fuel usage per flight). Others, such as buffer time, continue to increase. The average total cost per flight (the sum of fuel cost, CO_2 cost and airspace en-route charges) is very similar for the current, 2050 baseline and 2050 supportive case, while significantly higher for the 2050 non-supportive case. The main explanation is the sharp increment in fuel cost for the non-supportive scenario. Note in Table 50 how the MTOW of the aircraft increases with respect to the current scenario.

Table 50. Average MTOW and number of flight plan options 2050 supportive/non-supportive

Scenario	MTOW	Flight plan options
Current	85.7	8.9
2050 supportive	87.5	9.7

Scenario	MTOW	Flight plan options
2050 non-supportive	87.8	6.5

(ii) ANSPs

The revenues obtained by the ANSPs in the 2050 supportive and non-supportive scenarios are in the order of magnitude of the current case (see the following tables and figure). It is worth noting that the 2050 non-supportive case obtains less revenues than the current one, and that the 2050 supportive case is lower than for 2035 supportive (see Table 48).

Table 51. Total en-route airspace revenue M EUR 2050 supportive/non-supportive

Scenario	En-route airspace charges revenue M EUR
Current	26.2
2035 supportive	27.2
2050 non-supportive	22.4

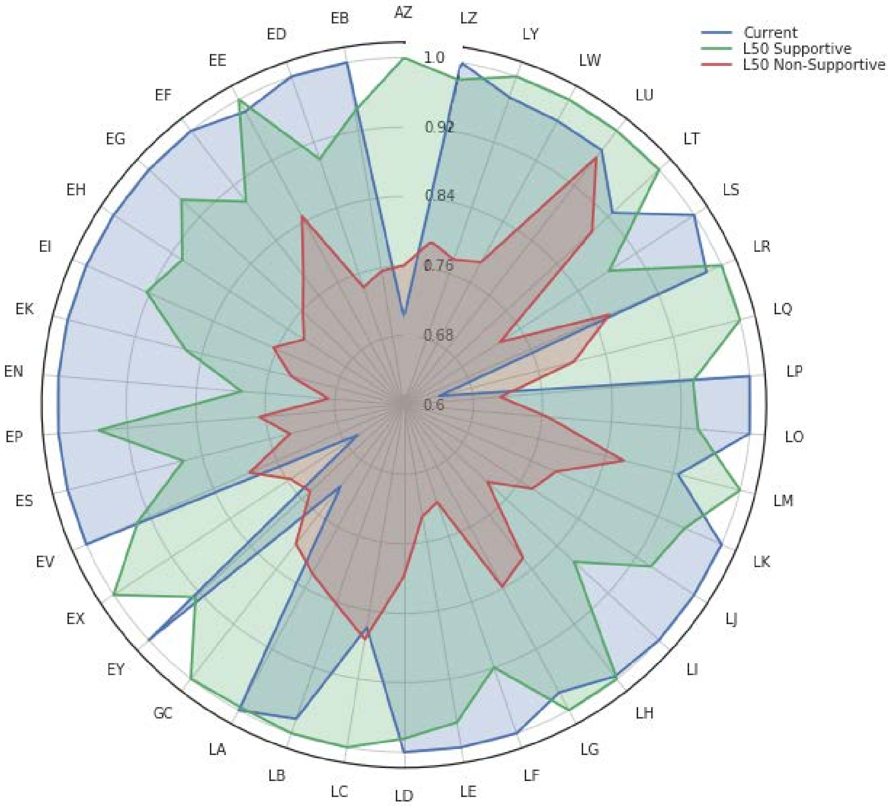

Figure 91. ANSP revenues due to en-route charges variation, 2050 supportive/non-supportive

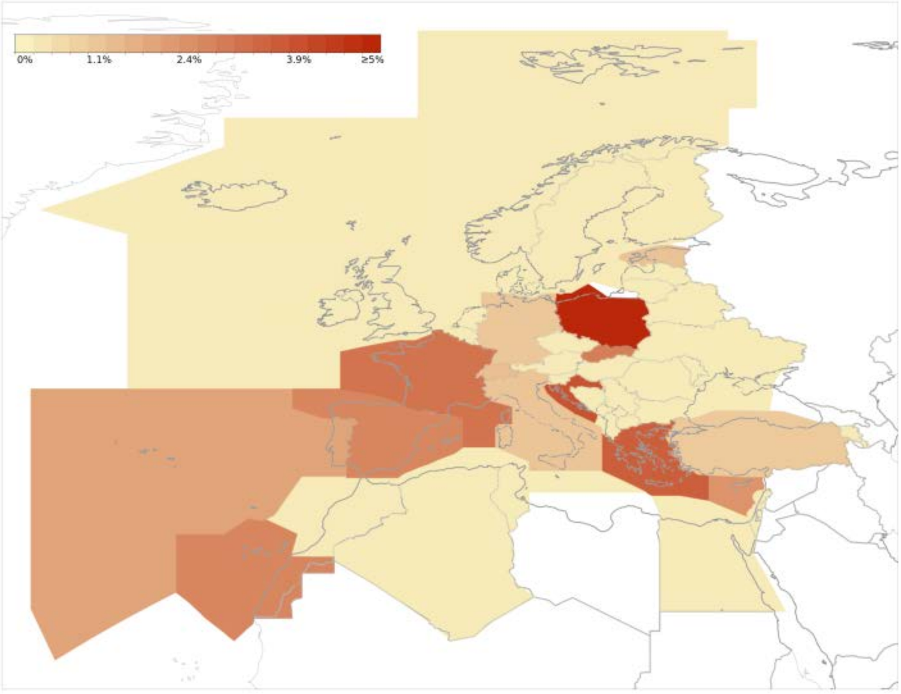

a) Current

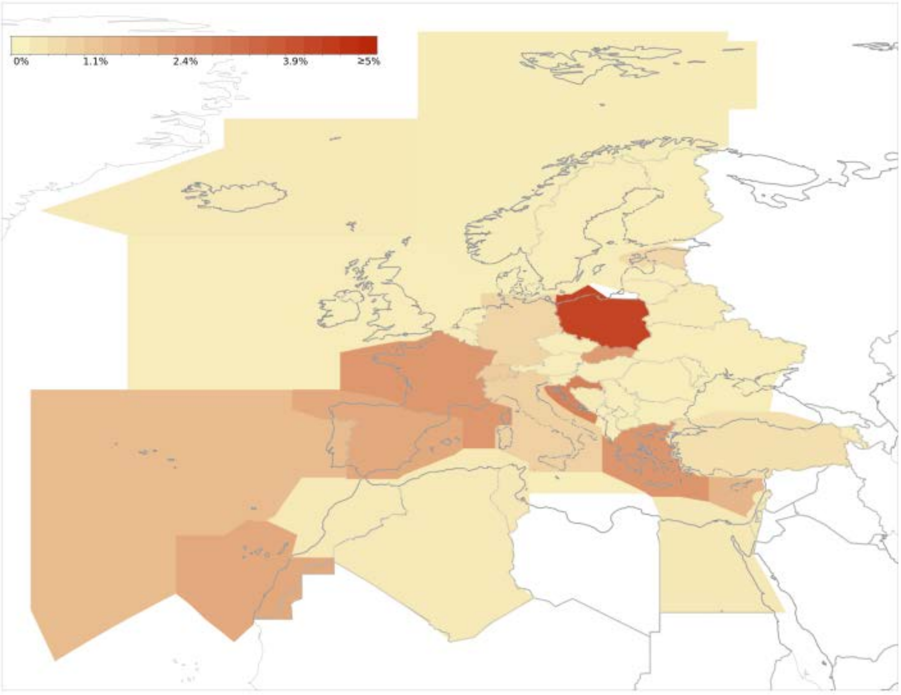

b) 2050 supportive

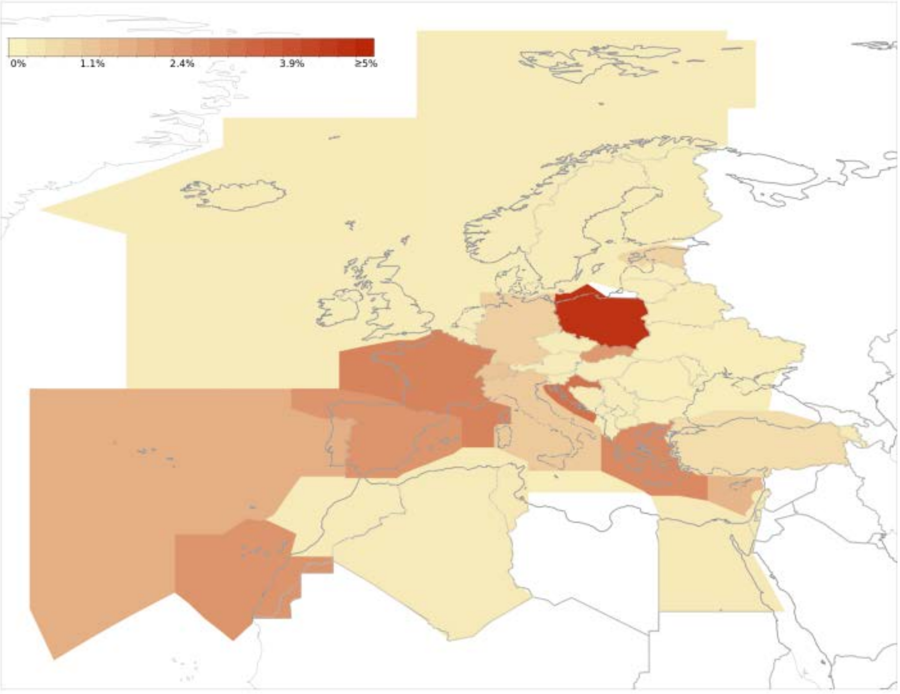

c) 2050 non-supportive

Figure 92. Probability ATFM delay capacity 2050 supportive/non-supportive

Regarding the probability of having delay due to ATFM capacity regulations, as with the 2035 cases, the probability tends to decrease with respect to the current scenario, being lower in the supportive than in the non-supportive scenario (see Figure 92).

(iii) Passengers

For the total number of passengers, the non-supportive 2050 case has a small increment (13.8%) relative to the current scenario, while in the supportive case, the increment is greater (36.0%), as presented in the following table. These increments are higher than for the 2035 scenarios but lower than the 2050 baseline case (see Table 43).

Table 52. Passenger details 2050 supportive/non-supportive

Scenario	Total number passengers ('000)	Passengers with 1 connection ('000)	Passengers with 2 connections ('000)	Total passengers with connections ('000)
Current	3 401	275 (8.1%)	9 (0.3%)	284 (8.4%)
2050 Supportive	4 626	406 (8.8%)	13 (0.3%)	418 (9.0%)
2050 Non-supportive	3 869	363 (9.4%)	12 (0.3%)	375 (9.7%)

5.2.3 Summary of the results and key points

5.2.3.1 Key metrics

In the pre-tactical layer, the key metrics presented capture the main characteristics of the flight plans, the ANSPs and the passengers, as shown below.

- For fight plan characteristics, to describe how the flight plans change due to the different scenarios:
  - average fuel per flight (kg);
  - average distance flown (NM);
  - average flight time (minutes);
  - buffer time per flight plan (minutes).
- Wider ATM system status: in order to identify how the system in general is evolving, the amount and characteristics of the traffic and passengers, and the impact on ANSPs' operations:
  - number of flights (count);
  - average MTOW of the aircraft used (tonnes);
  - number of passengers (count);
  - connecting passengers (percentage);
  - ANSPs' revenues (euros).

In Figure 93, we present the evolution of these key metrics for the flight plan characteristics. The average fuel burn per flight in 2050 is slightly lower than for the current scenario, and lower for the non-supportive scenario. The total amount of CO_2 emissions is proportional to the fuel consumption and hence follows the same trend. The flight plan distance increases for 2035 and 2050, being larger for the supportive scenarios. This indicates the selection of longer routes and the operation of longer origin-destination routes, as demand increases. Note that for both the supportive and non-supportive cases, the increase is lower than in the (low) baseline. The average flight plan time evolves proportionally to the increment in flight plan distance. However, buffer times (time allocated for taxi in, taxi out and buffer for the flight plan) increase over time and are higher for supportive scenarios. This is important, as it might have an impact on the amount of delay actually incurred by the flights tactically. Whilst longer routes are being used with longer flight times, fuel consumption is flat. Initial investigation of this effect suggests that it may be due to the assignment of more fuel-efficient aircraft types to the routes.

Figure 93. Key flight plan metrics evolution

If we consider the evolution of the system metrics in Figure 94, we see that the number of flights increases over time, even if lower for the supportive and non-supportive scenarios than for the low baseline. This is in accordance with the increment in the number of passengers. The average MTOW of the aircraft increases over time, with no difference between the supportive and non-supportive cases. An increment in the percentage of passengers making connections points towards more operations on longer routes. From an ANSP perspective, the revenues obtained due to en-route airspace charges of the supportive case are similar to the baseline case, which increase in 2035 with respect to the current value, and then decrease. In the non-supportive case, the revenues decrease as the capacity that needs to be provided is offset by the increment in flights.

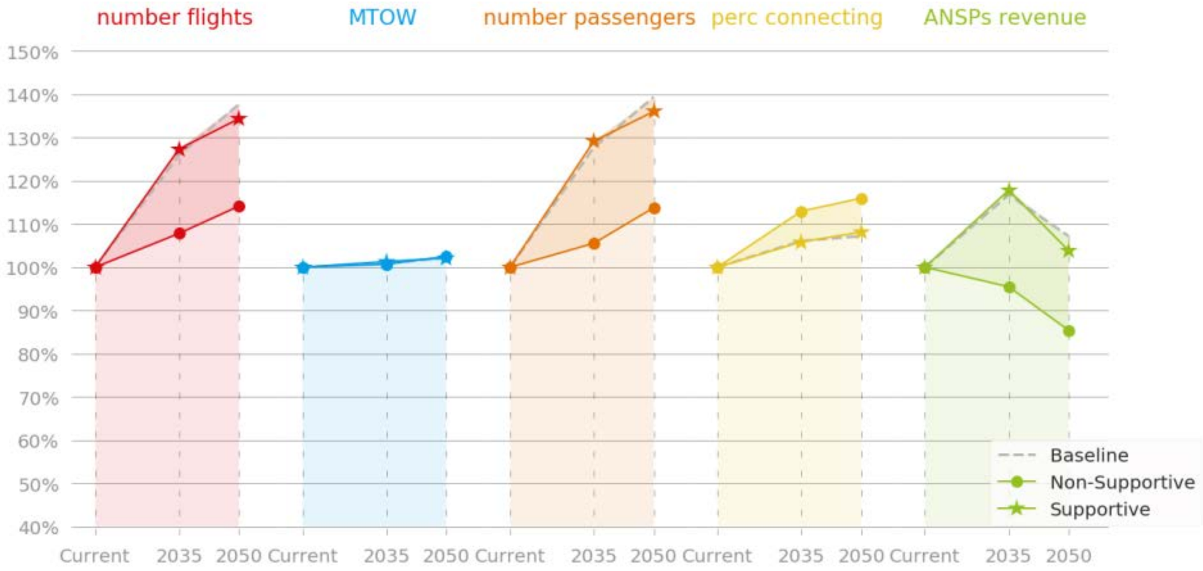

Figure 94. Key system metrics pre-tactical

5.2.3.2 Key points

We conclude this analysis by highlighting the main points that emerge from the pre-tactical results. These results obtained are in accordance with the outcome of the strategic layer.

1. Fuel consumption per flight is flat over time as the (e.g. technological) benefits obtained by the system are offset by the use of longer routes with larger aircraft, with a potential shift to greater fuel efficiencies. The relative importance of the fuel price over time might also favour the selection of trajectories that use less fuel.
2. The selection of larger aircraft over time is related to the increase of passenger demand and route length.
3. There is an increase in the size of the buffers per flight: this may contribute to the reductions in tactical delay costs (e.g. see Section 5.3) and could be used by the strategic layer to tighten the schedules – the reason for this increment may be linked to having more numerous longer routes, which usually have larger buffers to manage greater uncertainty.
4. The number of passengers connecting increases over time – these effects are also discussed further in Section 5.3.

5.3 Tactical layer

The Vista tactical layer reproduces the on-day operations from the flight plans produced by the pre-tactical layer and the passenger itineraries. The execution of flight plans may deviate from the planning phase for several reasons, such as weather, passengers missing connections, or flight cancellations. The metrics considered in the tactical layer are therefore performance oriented, in particular delays and the associated costs, *viz*.:

Flight-centric:

- departure delay;
- arrival delay;
- reactionary delay;
- gate-to-gate times;
- tactical delay costs;
- Fuel consumption;
- fuel costs;
- emissions costs.

Passenger-centric:

- departure delay;
- arrival delay;
- gate-to-gate times;
- door-to-door times.

5.3.1 Baseline evolution

The baseline scenarios reflect that in both cases the mid-term (2035) and long-term (2050) scenarios yield improved passenger delay metrics. There is a linear trend of improvement in both cases; the arrival delay almost halves. This is a very good result from the passenger perspective, as passengers are mostly sensitive to arrival delay. The improvement, however, has a diminishing return in the long term.

Moreover, gate-to-gate and door-to-door times are also reduced, although the total reduction is relatively marginal compared to the improvement in delay. This is a key finding, as it shows that the impact of flight delays on door-to-door times is fairly marginal and therefore reducing such delays might not be the best strategy to achieve a reduction in door-to-door times, i.e., the Flightpath 2050 vision of 90% of EU passengers travelling within four hours door-to-door for journeys involving air as one of the modes.

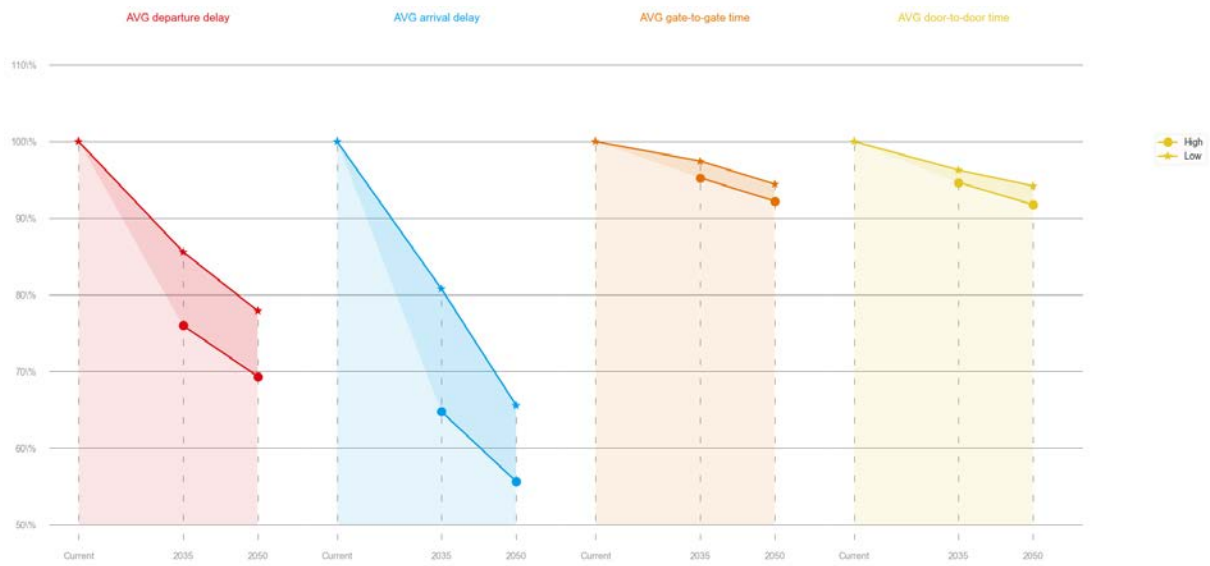

Figure 95. Passenger delay evolution on baseline scenarios

There is also an improvement in flight delay metrics. However, the relationships are not linearly correlated with the passenger metrics across all scenarios. The arrival delay for flights is reduced more than for passengers and it is more acute in the lower scenarios; high scenarios behave similarly in their mid- and long-term evolution. This is a fundamental fact that is usually overlooked: passenger delay and flight delays are related, but not the same. Therefore, reducing flight arrival delay has an impact on passenger arrival delay but the reduction factor will not necessarily be a constant linear relationship. Interestingly, the gate-to-gate time has a minimal increase, when compared to passenger gate-to-gate times, which are actually reduced.

Note that the passenger average gate-to-gate time is actually smaller than the average flight gate-to-gate time. Since some passengers have connections, one could expect the opposite to be true. However, this is a statistical artefact. Since a flight carries several passengers, the average gate-to-gate time for passengers is a weighted average of the gate-to-gate time for flights. In our model, shorter flights are more occupied and this drives the average value of passenger gate-to-gate times to be smaller. Also, connecting passengers usually take shorter flights, this effect compensates the fact that for connecting passengers gate-to-gate times are higher. Consider the extreme case of two flights and one passenger, if the passenger takes only the shorter flight, the gate-to-gate time for the passenger(s) would be smaller than the average gate-to-gate time of the two flights. This paradox is a special case of the amalgamation paradox that occurs when data are aggregated in a particular way, in this case passengers are aggregated by flights.

Regarding reactionary delay, in the mid- and long-term scenarios the percentage of flights propagating delay, and the average reactionary delay, decrease from the current scenario: this means fewer flights propagating even smaller delays. However, despite the reduction, the non-supportive mechanisms have a diminishing returns effect in the long term. The underlying mechanism for this is a saturated network with insufficient resilience.

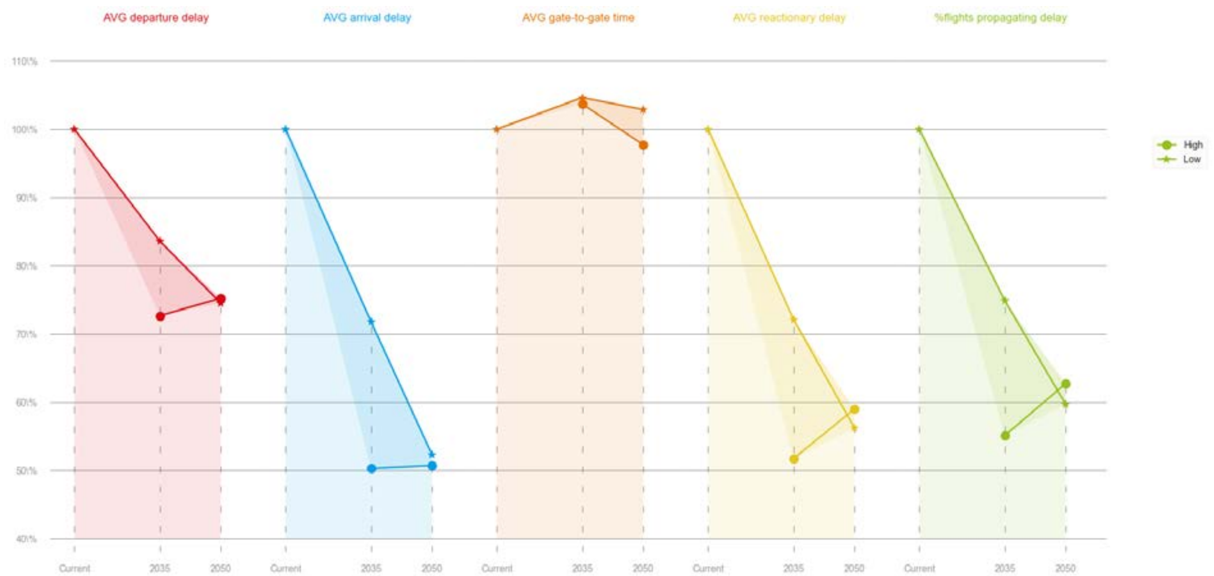

Figure 96. Flight delay metrics evolution in baseline scenarios

Tactical costs of delay are estimated considering the average arrival delay and MTOW, as described in [7]. These tactical costs of delay also decrease with the background improvements in the mid- and long-term. However, average fuel and emissions costs in the mid- and long-term remain relatively unchanged (5% with respect to the current scenario). There is a notable increase in the total fuel consumption. This is expected, and due to the increase in the number of operations. The higher fuel consumption leads to a total increment in the impact of aviation on the environment. The high scenario leads to an increase in both fuel costs and emissions costs.

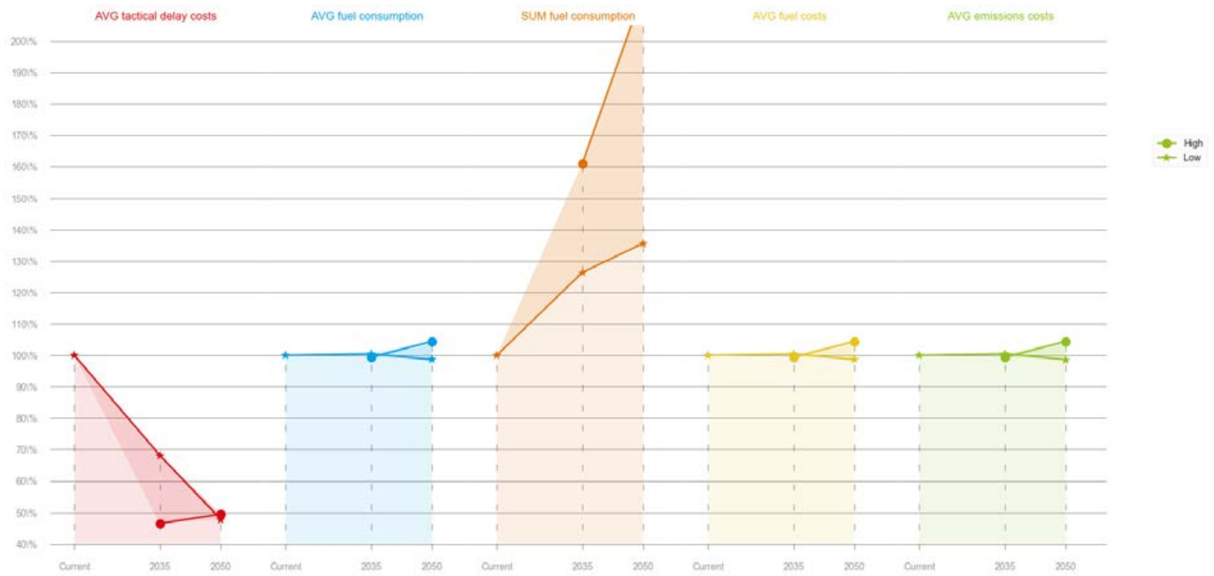

Figure 97. Flight cost metric evolution in the baseline scenarios

5.3.2 Supportive and non-supportive scenarios

As with the strategic and pre-tactical layers, the analysis of the supportive and non-supportive scenarios is done first for the 2035 cases, and then for 2050.

5.3.2.1 Supportive/non-supportive in 2035

In the mid-term (2035), the effect of both the supportive and non-supportive foreground factors has a positive impact on passenger delay metrics. The impact is, however, variable. The effect on departure delay is small in the non-supportive scenario, whilst it is quite large in the supportive case. The impact is more noticeable for arrival delay, and, similar to the baseline evolution described previously, the arrival delay is more reduced than departure delay. Total travel times, gate-to-gate and door-to-door, are also reduced in both scenarios. In this case, non-supportive factors lead to greater reductions. To find a satisfactory explanation, we have to look closely at the trade-off with the flight cost metrics in Figure 99.

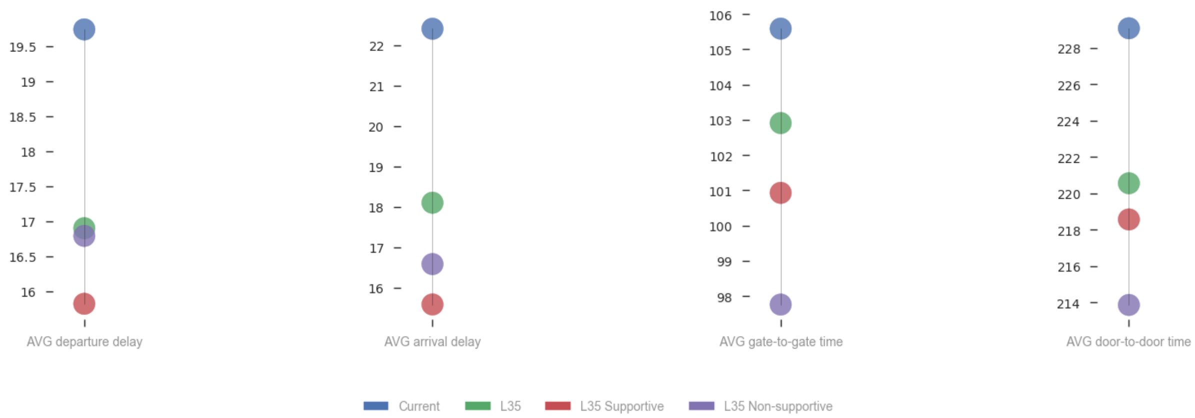

Figure 98. Passenger delay metrics for supportive/non-supportive scenarios in 2035

Note that the demand in the different scenarios may be different. Therefore, the slightly higher arrival delay of the non-supportive scenario with respect to the supportive one leads, on average, to a lower gate-to-gate time in the non-supportive case, indicating that, on average, flights are shorter.

Gate-to-gate times are further reduced in both scenarios, while reactionary delay and the number of flights propagating delay behaves oppositely in the supportive and non-supportive scenarios. This indicates how operators adapt to the imposed mechanisms using buffers and distributing delay across flights. Overall, as shown in Figure 99, the percentage of flights propagating delay decreases with respect to the current scenario, while the amount of delay propagated also tends to decrease.

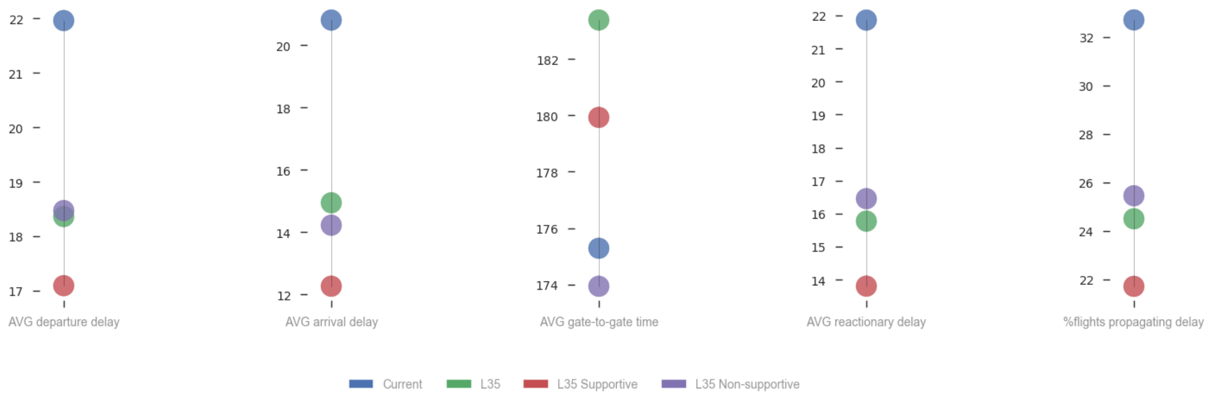

Figure 99. Flight delay metrics for supportive/non-supportive scenarios in 2035

One aspect that is often overlooked in purely performance-focused studies is the trade-offs with other metrics, in particular costs. Delay reductions are often achievable by increasing costs: whilst the above results look promising from a purely delay-performance point of view, the flight cost metrics show otherwise. Although tactical costs are indeed reduced in both the supportive and non-supportive scenarios, it is only at the expense of a large increase in fuel and emissions costs. Whilst the supportive scenario sees fuel costs at the same level as the baseline, the emissions costs increase in both scenarios. The impact of fuel cost is highly driven by the unit rate of fuel used in the scenarios and, therefore, in the supportive case, even if more fuel is used than in current operations, the cost is maintained at a similar level.

This trade-off in time performance versus total cost (tactical cost of delay plus emissions and fuel costs), is one of the main findings of the tactical layer. From the purely economic perspective, emissions and fuel costs can be partially compensated with a reduction in tactical costs of delay. From an environmental perspective, the average fuel consumption, and hence emissions per flight, is decreasing (even if not substantially), so the average impact per flight is lower than in the current scenario. Nevertheless, the number of flights in the scenario is increasing, so even if the impact of a flight on the environment on average is lower, the overall impact is higher than current simply because the traffic grows faster than the reduction in emissions. Note that Vista has not considered further improvements in aircraft engine fuel emissions reductions and/or the use of alternative fuels. The results show the importance of investing in those fields if the total impact of aviation aims to be maintained or reduced.

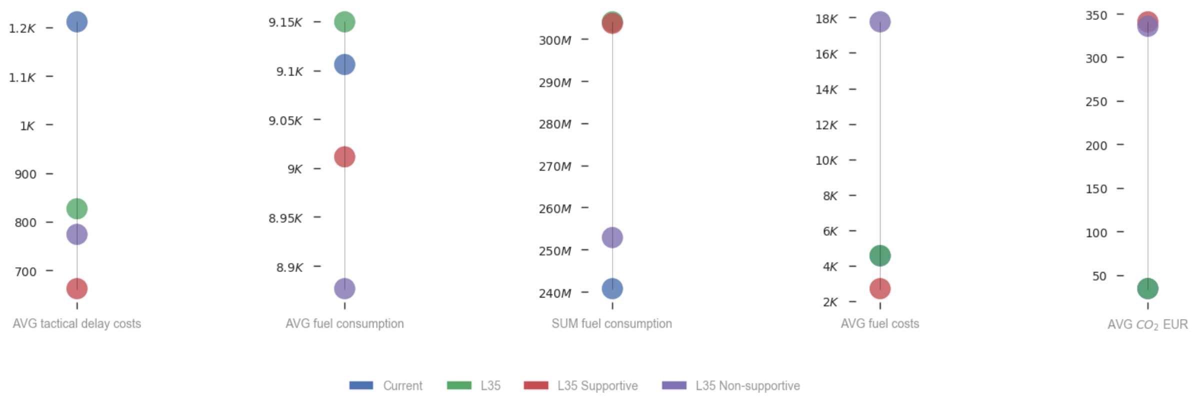

Figure 100. Flight cost metrics for supportive/non-supportive scenarios in 2035

5.3.2.2 Supportive/non-supportive in 2050

In the long-term (2050) scenario, the effects shown in the mid-term (2035) are more accentuated. Passenger departure and arrival delays are improved, although less considerably, and gate-to-gate and door-to-door times are also reduced overall.

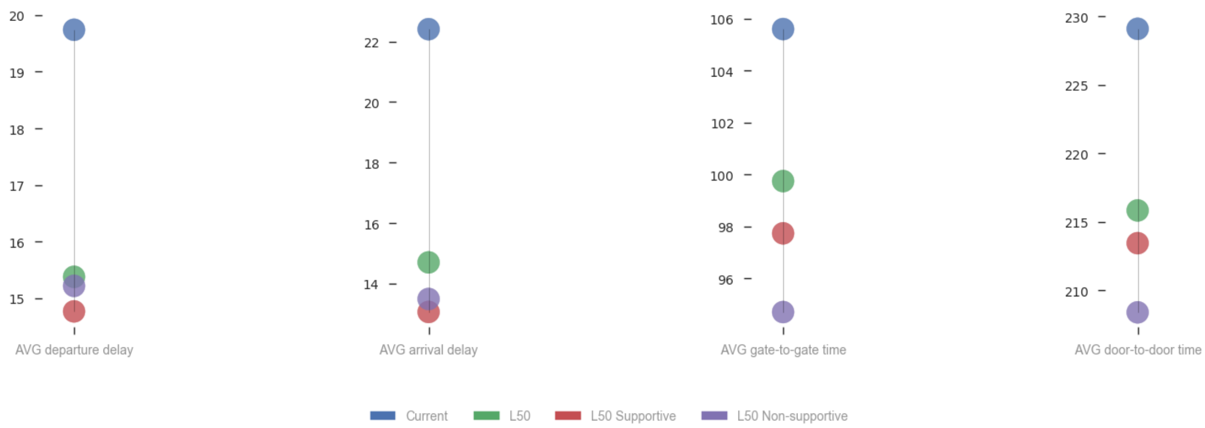

Figure 101. Passenger delay metrics for supportive/non-supportive scenarios in 2050

Flight metrics in the long-term show a similar trend as for the mid-term, at least for departure and arrival delay, and total gate-to-gate times. Reactionary delay, however, behaves slightly differently. Both the supportive and non-supportive scenarios remain close in 2050, whereas there is an inversion with respect to the results of 2035. The difference is, however, less than 5%, which could indicate that this is just an artefact, or simply that the system oscillates trying to find a stable configuration. Either way, further studies will be needed to find the cause.

There is also a diminishing returns effect with respect to the number of flights affected by the reactionary delay in the long-term. The positive effects of the supporting factors are reduced over time. This is a common effect of new technologies and developments, which is typically difficult to quantify, but which the Vista tactical layer could help to quantify in future modelling.

Figure 102. Flight delay for supportive/non-supportive scenarios in 2050

In the long-term scenario, the trade-off between on-time performance and costs is again confirmed. Practically the same results (on a different scale) are obtained. Tactical costs are reduced further, although not as much as in the mid-term scenario, while emissions costs increase. Fuel costs are reduced in the supportive scenario and increased in the non-supportive scenario. However, when looking at the total fuel consumption, it is decreased further in the supportive scenario, even with the increase of total number of flights, than in the non-supportive. Interestingly, when looking at average consumption, this is reduced similarly in both scenarios.

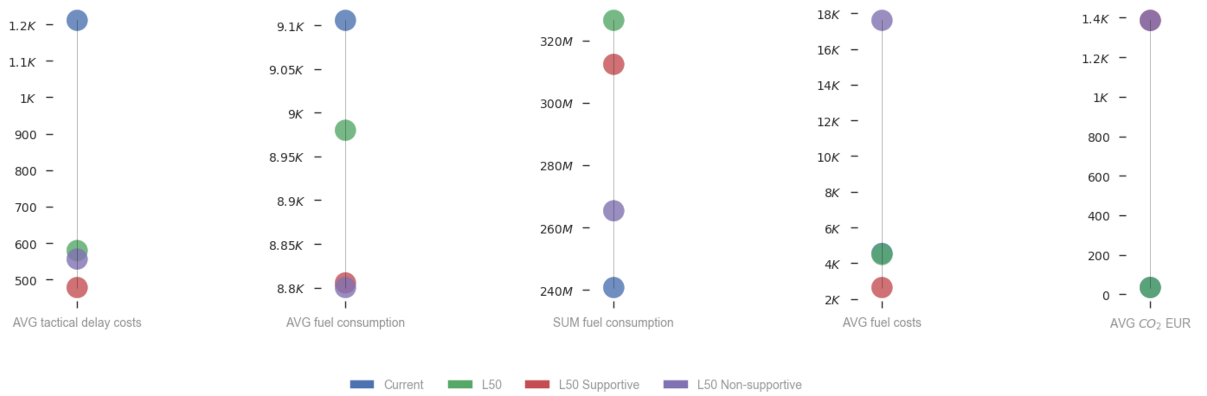

Figure 103. Flight costs for supportive/non-supportive scenarios in 2050

5.3.3 Summary of the results and key points

5.3.3.1 Key metrics

The key metrics for the summary have been selected to highlight the differences between passenger and flight centred metrics.

- Passenger-centred:
  - Departure delay;
  - Arrival delay;
  - Gate-to-gate time;
  - Door-to-door time.
- Flight-centred:
  - Departure delay;
  - Arrival delay;
  - Gate-to-gate time;
  - Reactionary delay;
  - Percentage of flights propagating reactionary delay.

When comparing the passenger metrics with the flight metrics, there are clearly similar trends in both departure and arrival delay. Whilst the relationship is clear, the metrics do not scale the same way. This result is repeated across all scenarios and simulations. The reduction in arrival delay on average is quite substantial: close to 50%. There are several rationales for this phenomenon.

Since the only relationships between departure and arrival delay is the airborne phase of the flight, the reduction in delay may be due to any of the following: shorter routes than expected tactically; better winds than expected; buffer times which absorb the departure delay and thus not propagating it into arrival delay.

In the tactical model, routes are selected as produced by the pre-tactical layer and winds are modelled statistically so that the average impact should be null. Departure delay, however, decreases in line with the reduction of reactionary delay, i.e. delay is not being propagated as much. In general, longer flights have larger buffers planned. The fact that in the future we have longer flights (which can be seen in the increase of the gate-to-gate time) leads to more buffer in the schedules (as they are based on historical schedules). However, the routes are shortened (e.g. due to free-routing) leading to larger buffers relative to the fixed schedules. The system might also be more predictable such that those buffers are not as necessary. This is an example of a feedback loop required to tighten the schedules strategically according to tactical operations.

It is also striking to observe that even if gate-to-gate times increase for flights in the mid- and long-term scenarios, it is in fact reduced from the passenger perspective. Despite gate-to-gate times for passengers being composed of one or more flight and gate-to-gate sequences, there are also the layovers for connecting passengers. Reducing passenger waiting times at the airport makes the overall trip length shorter. However, lower connecting times could lead to reduced passenger spend, which in turn could increase airport charges and ultimately passenger ticket prices.

Finally, from a flight perspective, the average amount of delay that is propagated as reactionary delay decreases over time, and the percentage of flights which propagate delay also decreases. This, once again, supports the idea that large buffers are used in the scenarios and delay is absorbed leading to lower arrival delay and hence propagation of reactionary delay.

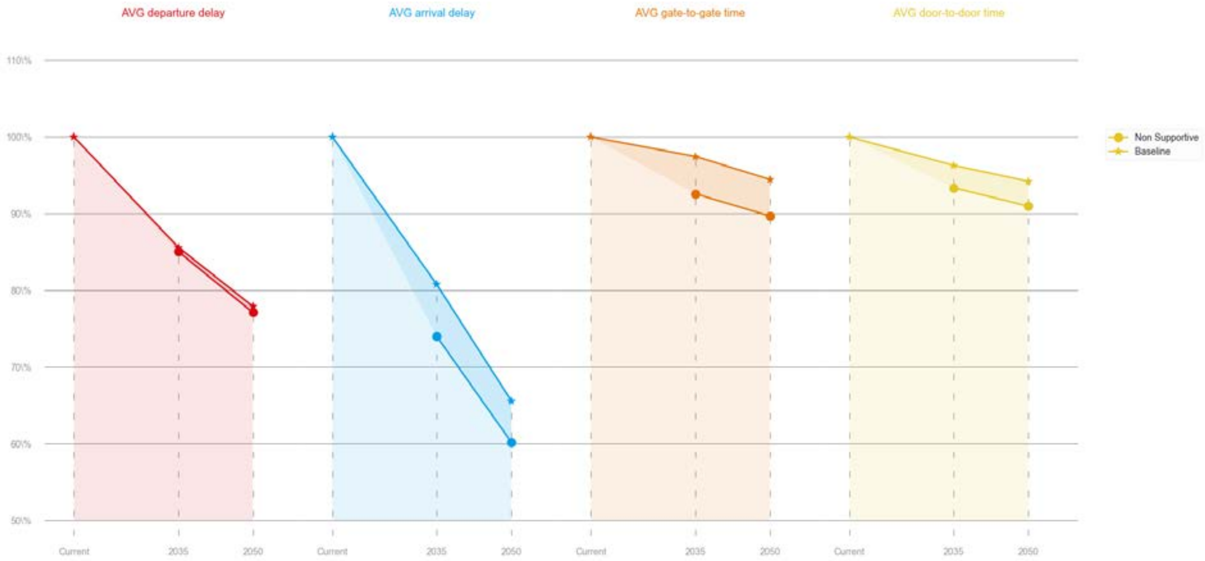

(a) non-supportive

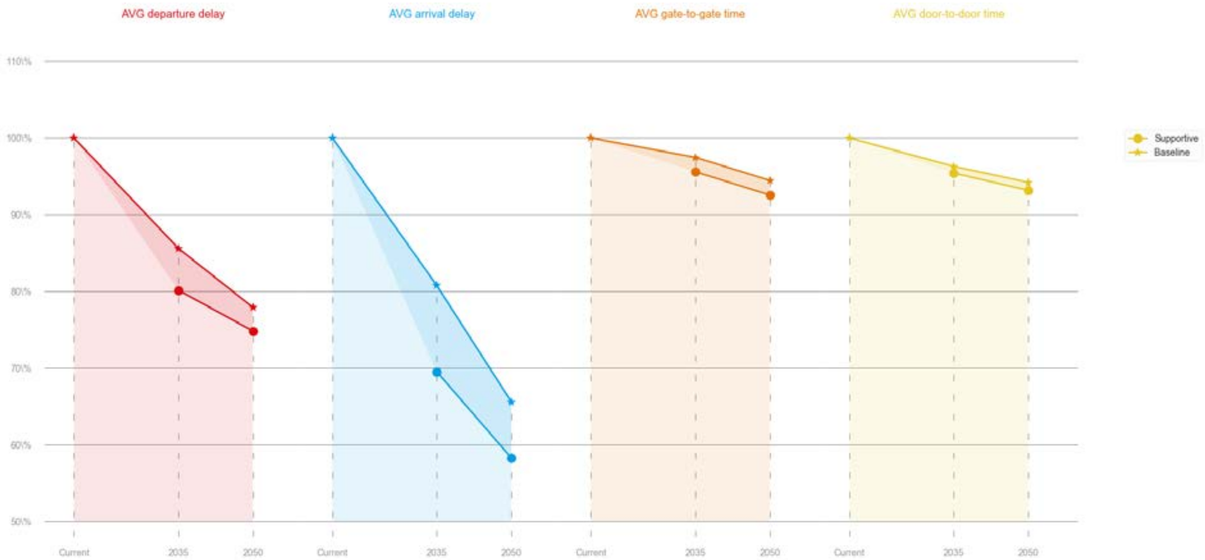

(b) supportive

Figure 104. Evolution of passenger delay metrics in supportive and non-supportive scenarios

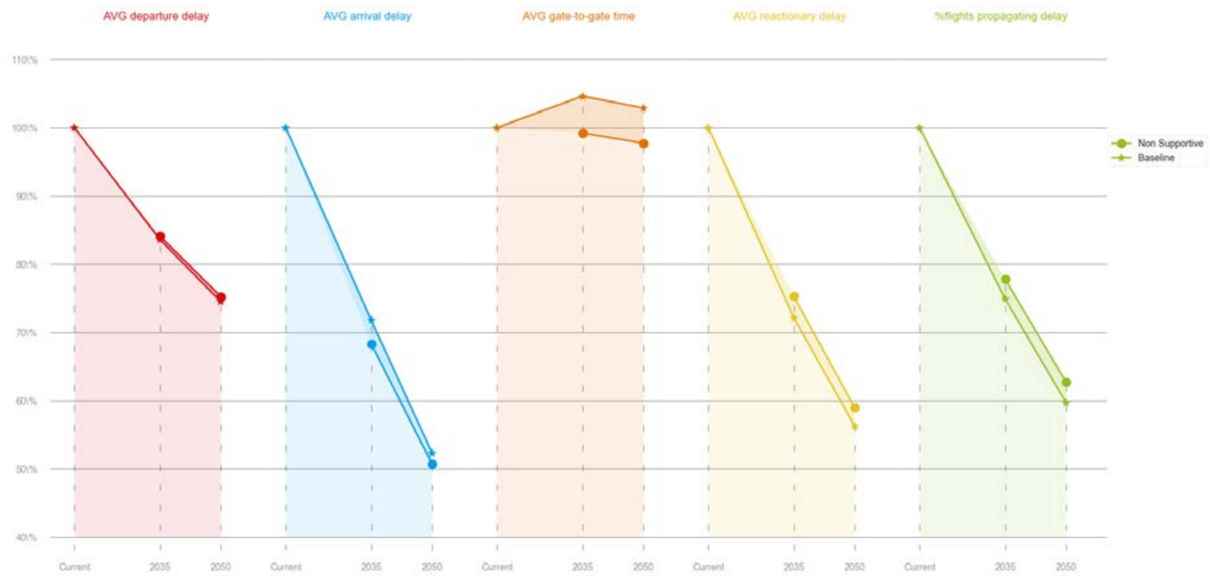

(a) non-supportive

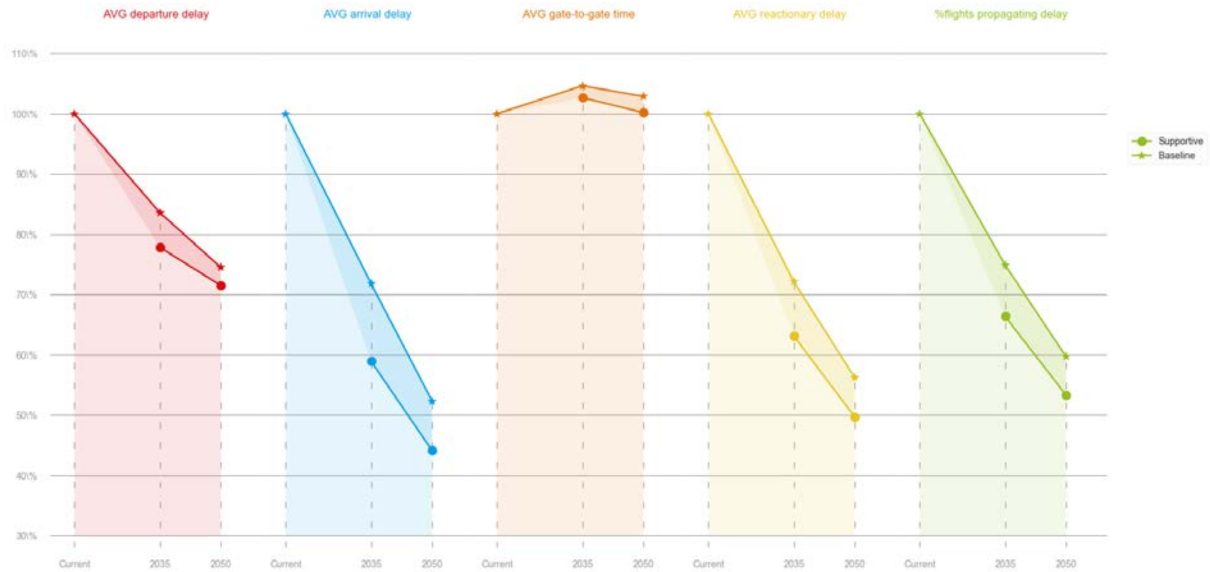

(b) supportive

Figure 105. Evolution of flight metrics in supportive and non-supportive scenarios

5.3.3.2 Key points

To summarise, the main findings of the tactical layer in Vista are as follows.

1. Most passenger- and flight-centric metrics follow similar trends overall, but there are non-linear differences on how metrics scale, i.e., no simple, direct translation between passenger and flight metrics.
2. Reducing delay, either departure or arrival, has a limited effect on the total door-to-door travel times for passengers.
3. Reductions in flight arrival delay with passenger arrival delay map close to a 1 : 1.3 ratio.
4. There is a diminishing return of the positive effects of supportive factors (mechanisms) in the long-term (e.g. 2050).
5. There is a clear trade-off between delay performance and cost metrics: improving system performance is usually expensive; Vista can quantify such trade-offs to find compromise solutions.
6. The results show that an improvement in passenger door-to-door times does not necessarily imply an increase in the average emissions per flight.
7. The average emissions per flight tend to decrease over time but the increase in traffic leads to an overall higher impact of aviation on the environment.

6 Future work and next steps

6.1 Model capabilities and future work

Vista is more than a simple model. It is an integrated set of three models: strategic, pre-tactical and tactical. Each is composed of different blocks and can be executed independently or jointly, the outcome of one layer being the potential input to another. This allows us to perform analyses not only such as those presented in this deliverable, but also targeted case studies such as those described in D5.1 (where the impact of a given variable is analysed more finely by each of the layers, independently).

The modularity of the development of the model layers allows the enhancement or part replacement of the model seamlessly. Vista is capable of capturing and quantifying the relationship between complex metrics across several stakeholders, reproducing classic KPIs and estimating complex and newly defined ones (such as the cost of uncertainty, or door-to-door travel times for passengers).

This modularity implies that the different parts can be adapted and reused in other situations. The strategic and pre-tactical layers will, for instance, be used in current and future projects to produce synthetic data to be used by other models, such as Domino and ADAPT in SESAR ER3.

The simulations already performed in Vista have created a large database, which contains many metrics that can be analysed more in depth, e.g. regarding in detail the ATFM delay generated per ANSP for different scenarios, or variations in specific flows of passengers through Europe. This includes deeper analyses of the behaviour of the model in different situations and with different parameters, furnishing further insights into future scenarios and trade-offs.

Vista is also unique in the sense that it supports the analysis of how a given metric changes during the temporal evolution of the different ATM phases: from the expected outcomes of the stakeholders' plans defined strategically, to the planned operations pre-tactically, to the actual execution phase, tactically. As future work, the outcome of the downstream layers can be fed back to the previous layers to improve how some decisions are made. This is a typical case of reinforcement learning, which can be used to make better predictions and also to optimise the system.

Further scientific questions can be studied by modifying the model. The impact of infrastructure expansions for airports, the way airlines may compete for new routes, and new ANSP structures (including pricing schemes) are but some examples. Further internal mechanisms (new behaviours) and external factors could be added.

Concerning the latter, a comprehensive list was compiled in D3.1. Some of them would require little effort and only moderate changes to the model. Others would need more dedicated research and development, for instance concerning the impact of drones on the ATM system.

Finally, Vista has developed several subroutines for an automated analysis of the results, including statistical tests and graphical representations. Given the amount of information produced by the model, a more interactive interface would be very useful to explore the data. This is a common issue in this kind of project, where the effort is mainly focused on the development and production of results, and cannot expend effort on producing interfaces for non-experts.

6.2 Next steps

The next steps in Vista focus on maximising the dissemination and increasing the opportunities for further development of the tool and its usage by different stakeholders. The consortium is in contact with several institutions that have shown interest in the capabilities of a system such as Vista.

These include discussion with PJ.19, facilitated by the SJU, with particular regard to the trade-off capabilities of Vista, and also with the European Commission regarding policy evaluation (impacts on door-to-door times from changes in Regulation 261, in particular).

Further activities with the airline and ANSP members of the consortium to explore how Vista could help them to assess the impacts of strategic decision-making are also foreseen.

A number of meetings with commercial entities, service providers and governmental bodies are also already planned. These are not reported here for reasons of confidentiality.

It is also planned that the members of the consortium will submit several articles to peer-reviewed journals and participate in conferences and workshops where Vista's approach and/or results are apposite. This will be subject to the availability of alternative funds for these activities.

7 References

1. COMMISSION REGULATION (EU) No 601/2012 of 21 June 2012 on the monitoring and reporting of greenhouse gas emissions pursuant to Directive 2003/87/EC of the European Parliament and of the Council.
2. EUROCONTROL, European Aviation in 2040 – Challenges of Growth, 2018.
3. Gurtner, G., Cook, A., Graham, A., Cristobal, C, “The economic value of additional airport departure capacity”, Journal of Air Transport Management, 69, 1-14, 2018.
4. IATA, “Air travel demand”, IATA Economics Briefing no 9, 2008.
5. Adler, N., “Competition in a deregulated air transportation market”, European Journal of Operational Research 129, 337-345, 2001.
6. DATASET2050 H2020 project, Deliverable 3.1, “Current passenger profile”, 2017.
7. Cook, A., Tanner, G. “European Airline Delay Cost Reference Values, Updated and Extended Values”, version 4.1, 2015.
8. Performance Review Commission, “Performance Review Report 2016”, June 2017.
9. EUROCONTROL, CODA DIGEST - All-Causes Delay and Cancellations to Air Transport in Europe – Annual 2014, 2015.
10. STATFOR, “Challenges of Growth”, 2013.
11. EUROCONTROL, Monthly follow up of service Units All CRCO Charging zones, 2018.
12. EUROCONTROL, Monthly adjusted unit rates, February 2018.
13. ICAO, “ICAO DATA+”, http://www4.icao.int/newdataplus, accessed September 2017.
14. Dahlmann, K., Koch, A., Linke, F., Lührs, B., Grewe, V., Otten, T., Seider, D., Gollnick, V., Schumann, U. “Climate-compatible air transport system – climate impact mitigation potential for actual and future aircraft”, Aerospace, 3(4), 38, 2016. doi:10.3390/aerospace30400338
15. European Energy Exchange, https://www.eex.com/en/market-data.
16. Global CCS Institute, Carbon Valuation in UK policy appraisal: a revised approach, July 2009.
17. P. Luckow and E. A. Stanton and S. Fields and B. Biewald and S. Jackson and J. Fisher and R. Wilson, 2015 Carbon Dioxide Price Forecast , Synapse Energy Economics Inc., March 3, 2015.

18. EUROCONTROL, Reporting Assumptions and Descriptions, Edition 6.25, November 6, 2017.
19. SESAR Joint Undertaking, ComplexityCosts D4.5 – Final Technical Report, October, 2016.
20. Regulation (EC) No 261/2004 of the European Parliament and of the Council. Establishing common rules on compensation and assistance to passengers in the event of denied boarding and of cancellation or long delay of flights, and repealing Regulation (EEC) No 295/91, 17 February 2004, 1-7.
21. European Parliament votes on air passenger rights, Press Release IP/14/119, Brussels, 5 February 2014, 1-4.
22. Cook A and Tanner G, 2011. A quantitative exploration of flight prioritisation principles, using new delay costs.Journal of Aerospace Operations, 1, 195-211.
23. Cook A, Tanner G and Lawes A, 2012. The hidden cost of airline unpunctuality.Journal of Transport Economics and Policy, 46 (2), 157-173.
24. EUROCONTROL, CODA Digest: All-Causes Delay and Cancellations to Air Transport in Europe – 2014. Ed. CDA_2014_005
25. Kano N., Seraku N., Takahashi F. and Tsuji S., 1984. Attractive quality and must-be quality, Hinshitsu, 14(2), 39–48.

8 Annex: Interpolated and extrapolated targets

Table 53. Time-based operations, SESAR targets used for factors

SESAR subpackage	Airport Cap.	Airspace Cap. (ER)	Airspace Cap. (TMA)	Cost effectiveness (ATCO)	Cost effectiveness (TECH)	Env./Fuel Eff.
01.01						
01.02						
01.03	7%					-0.25%
02.01		5%	5%	3.33%		-0.5%
03.01		2%	2%	0.5%		-0.65%
03.02	1%		5%	0.73%		-0.3%
03.03		8.5%	5%	5.83%		-0.21%
03.04						
04.01	1%		5%	0.73%		-0.3%
04.02						-0.25%
05.01						-0.02%
05.03		10%	4.65%	4.17%		-0.27%
06.01				0.84%	-1.55%	
06.03					-1.55%	
ENB01					-2%	
ENB02						-0.1%

Table 54. Trajectory-based operations, SESAR targets used for factors

SESAR subpackage	Airport Cap.	Airspace Cap. (ER)	Airspace Cap. (TMA)	Cost effectiveness	Flight Duration var.	Env./Fuel Eff.
01.01						
01.02						
01.03	2.14%		4.31%		-2.56%	-0.37%
02.01		7.65%			-1.71%	-0.69%
03.01	0.32%	3.06%	5.44%		-3.42%	-0.57%
03.02	1.07%				-4.70%	
03.03		13%	21.8%		-2.73%	-0.85%
03.04						
04.01	1.07%		6.74%		-4.70%	-0.05%
04.02	0.12%				-5.98%	-0.09%
05.01					-4.27%	
05.03		15.3%	4.4%			-0.22%
06.01	1.36%			11.85%		
06.03				11.85%		
ENB01				15.29%		
ENB02					-0.43%	

Table 55. Performance-based operations, SESAR targets used for factors

SESAR subpackage	Airport Cap.	Airspace Cap. (ER)	Airspace Cap. (TMA)	Cost effectiveness	Env./Fuel Eff.	Flight Duration var.
01.01						
01.02						
01.03	2.33%				-0.25%	-2.69%
02.01		19.98%	18.94%		-0.5%	-1.79%
03.01		7.99%	7.58%		-0.65%	-3.59%
03.02	0.33%		18.94%		-0.3%	-4.93%
03.03		33.97%	18.94%		-0.21%	-2.87%
03.04						
04.01	0.33%		18.94%		-0.3%	-4.93%
04.02					-0.25%	-6.27%
05.01			0.94%		-0.02%	-4.48%
05.03		39.96%	17.61%		-0.27	
06.01				11.85%		
06.03				11.85%		
ENB01				15.29%		
ENB02					-0.1%	-0.45%

9 Annex: full results

In this annex we present the main plots generated from the model. They are presented for reference, comparison, and completeness.

9.1 Economic model

9.1.1 Maps

Figure 106. Current scenario

Figure 107. H35 baseline scenarios

Figure 108. H50 baseline scenario

Figure 109. L35 baseline scenario

Figure 110. L35 Non-supportive scenario

Figure 111. L35 Supportive scenario

Figure 112. L50 baseline scenario

Figure 113. L50 Non-supportive scenario

Figure 114. L50 Supportive scenario

9.1.2 Key metrics plots

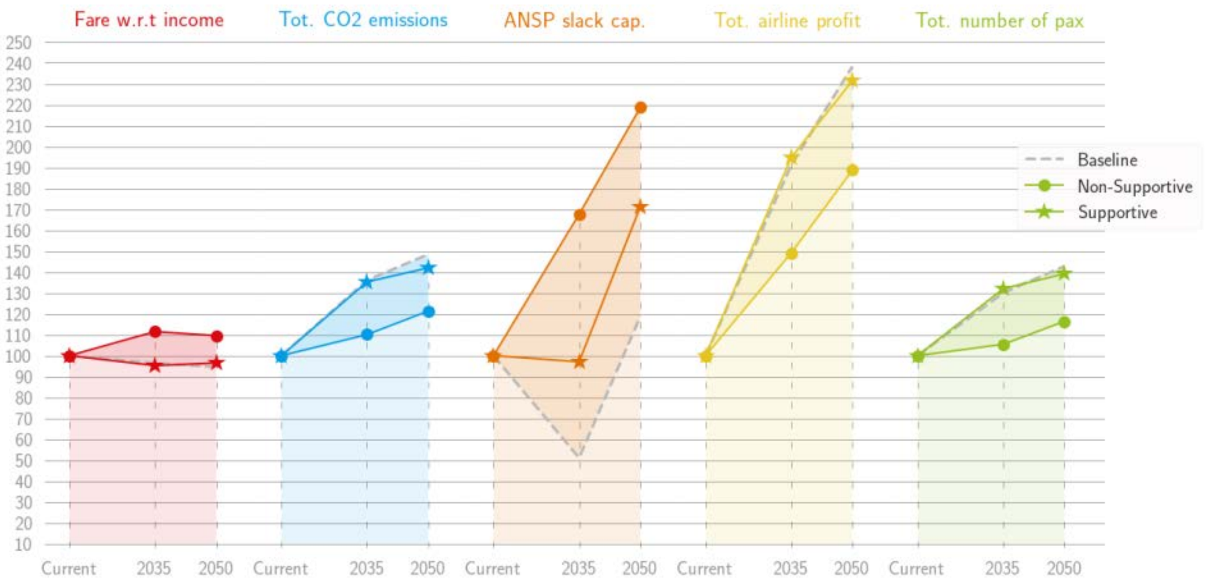
Fare w.r.t income
Tot. CO2 emissions
ANSP slack cap.
Tot. airline profit
Tot. number of pax
Baseline
Non-Supportive
Supportive
Current
2035
2050


Uncertainty
CO2 per flight
Departure delay
Cost unc. per flight
ANSP unit rate
Baseline
Non-Supportive
Supportive
Current
2035
2050

9.1.3 Line plots

9.1.3.1 Baselines

9.1.3.1.1 Low baseline

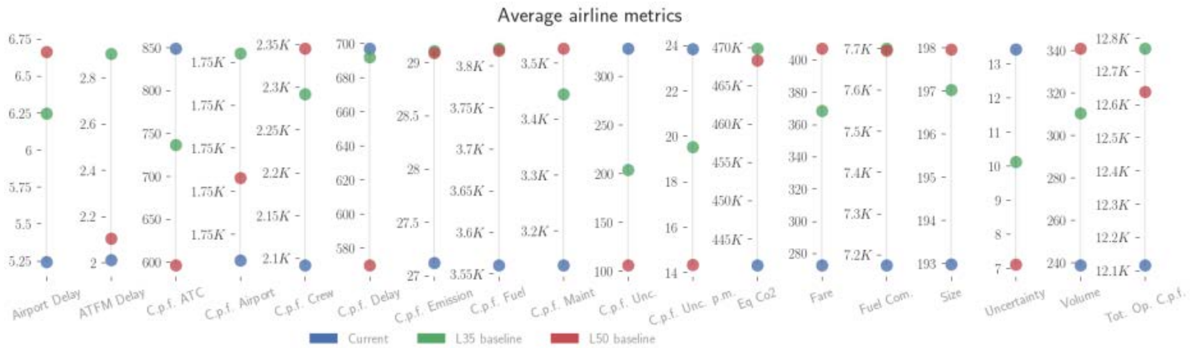

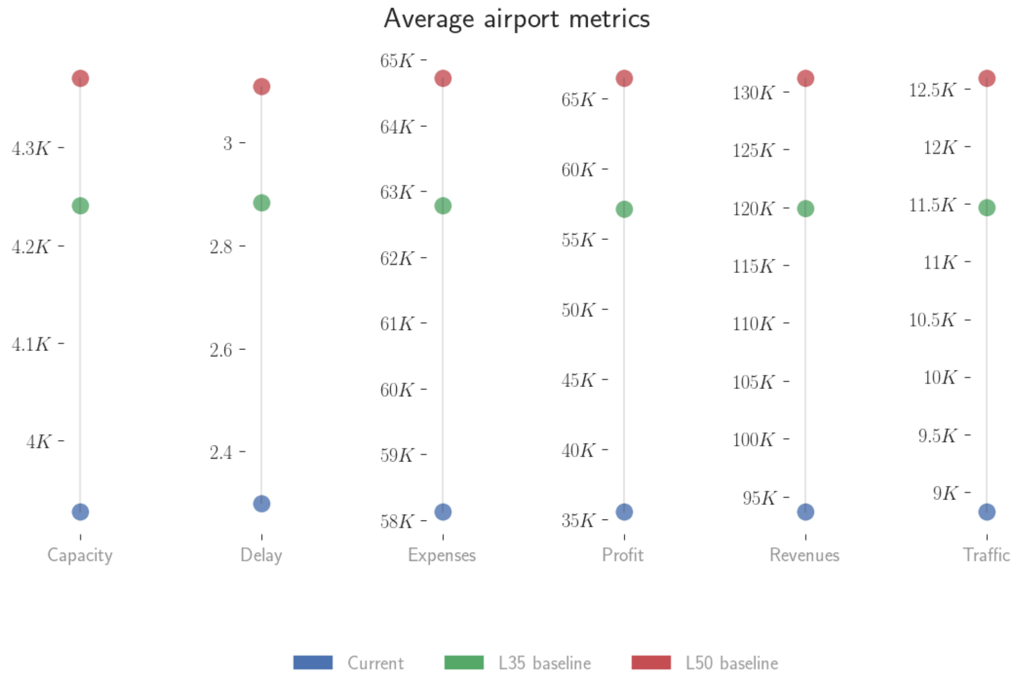

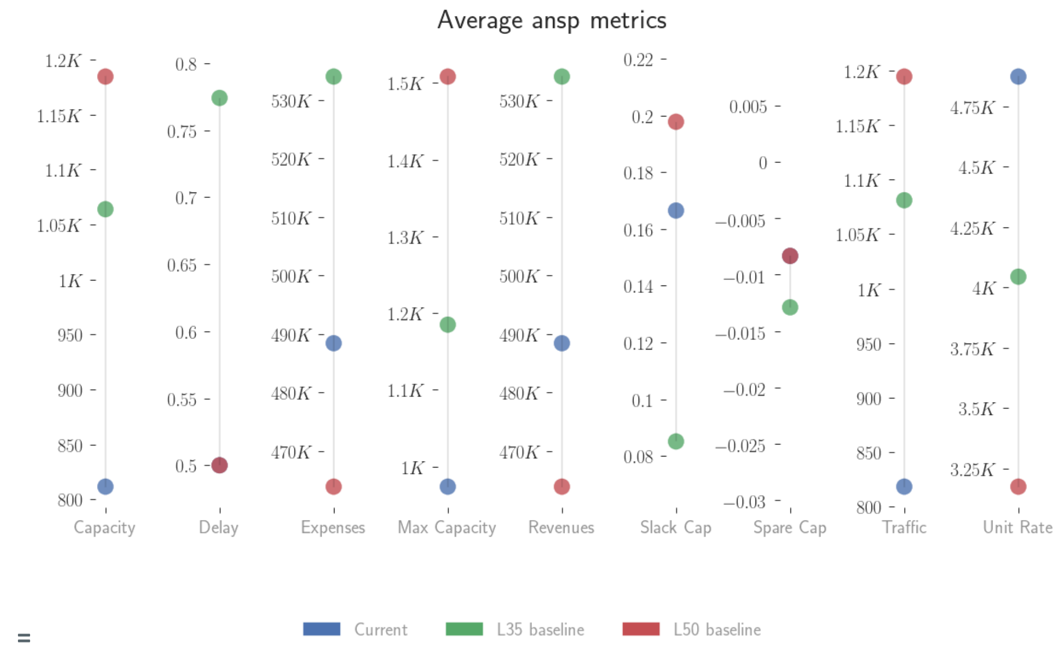
Average ansp metrics
Capacity
Delay
Expenses
Max Capacity
Revenues
Slack Cap
Spare Cap
Traffic
Unit Rate
Current
L35 baseline
L50 baseline

=

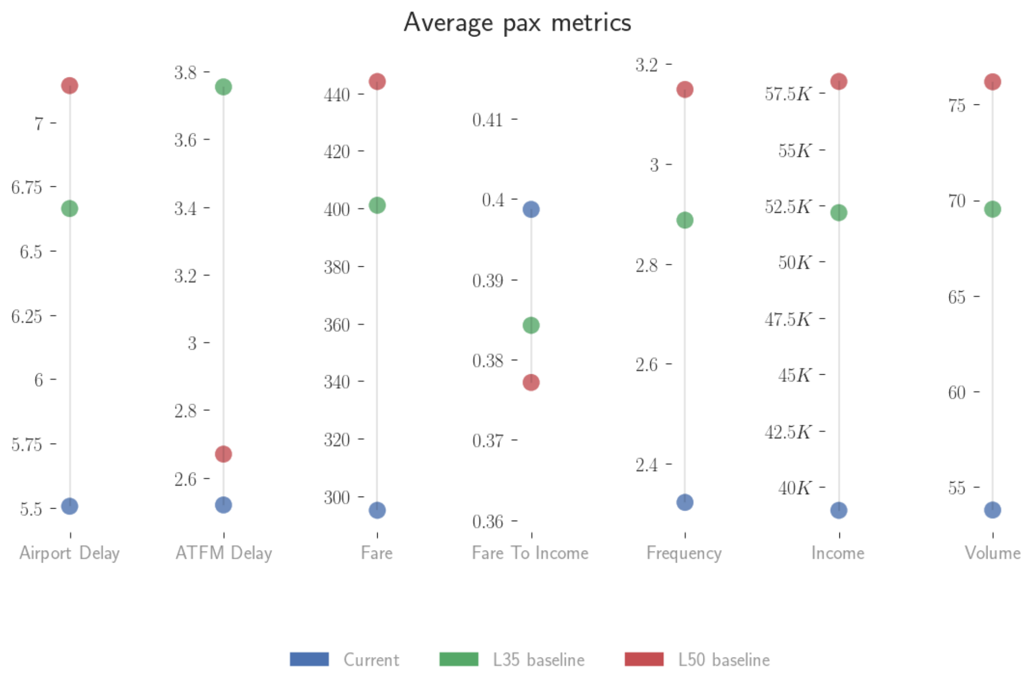
Average pax metrics
Airport Delay
ATFM Delay
Fare
Fare To Income
Frequency
Income
Volume
Current
L35 baseline
L50 baseline

9.1.3.1.2 High baseline

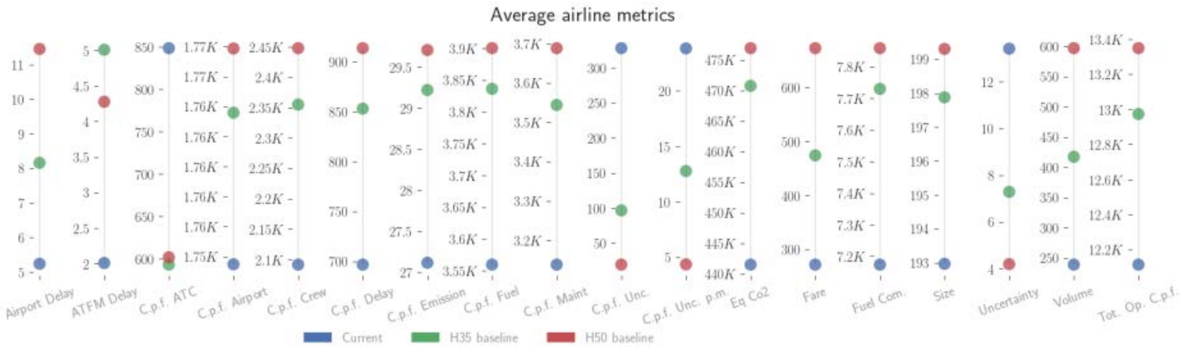

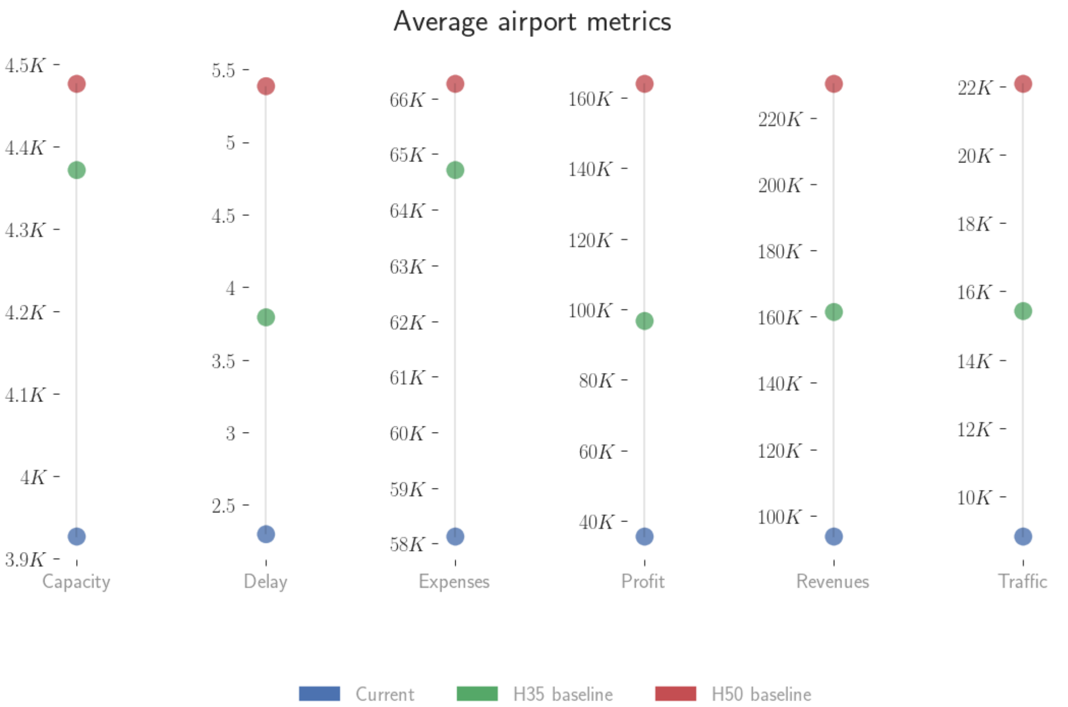

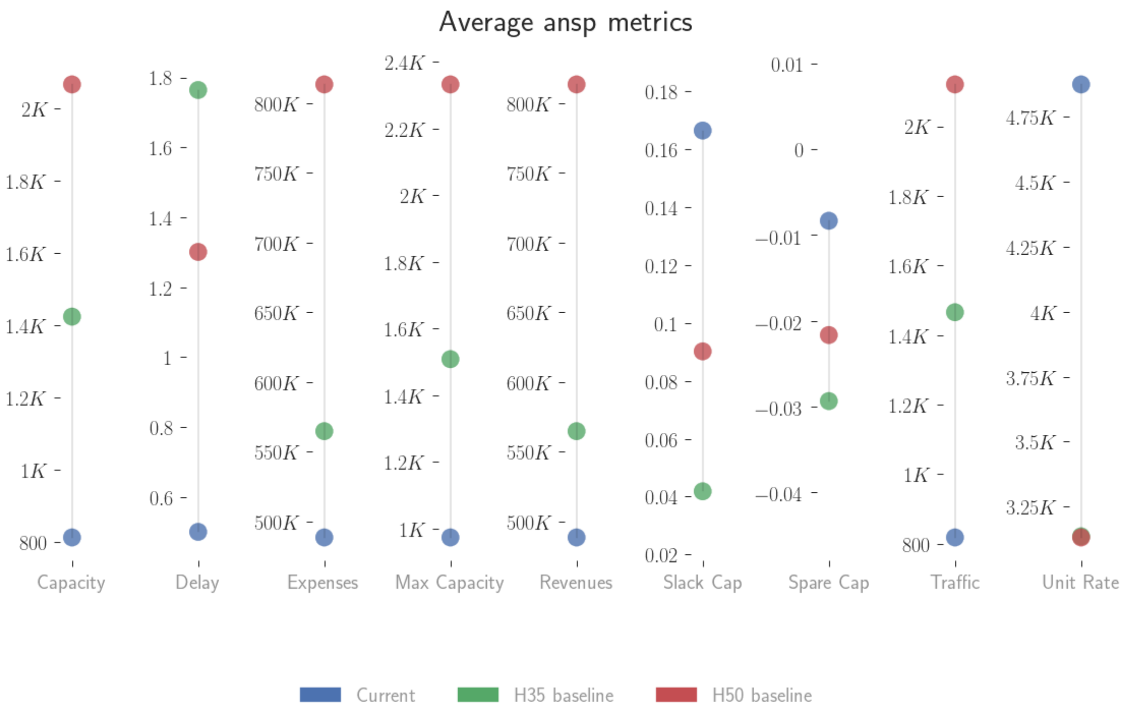
Average ansp metrics
Capacity
Delay
Expenses
Max Capacity
Revenues
Slack Cap
Spare Cap
Traffic
Unit Rate
Current
H35 baseline
H50 baseline

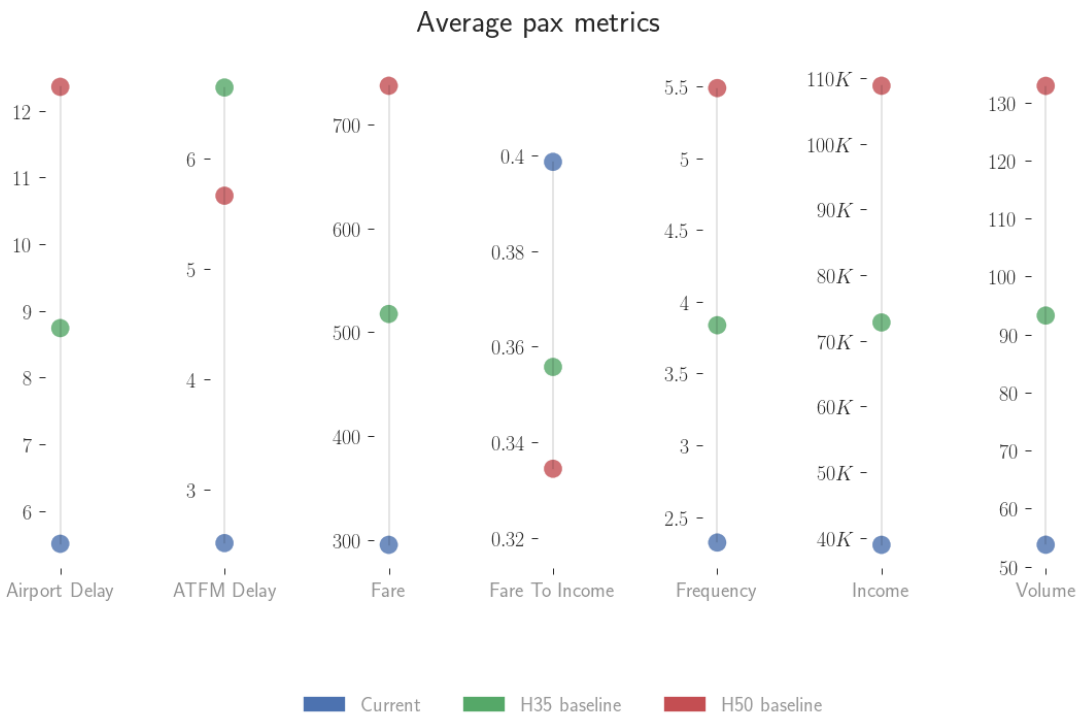
Average pax metrics
Airport Delay
ATFM Delay
Fare
Fare To Income
Frequency
Income
Volume
Current
H35 baseline
H50 baseline

9.1.3.2 Supportive/Non-supportive

9.1.3.2.1 Supportive/Non-supportive in 2035

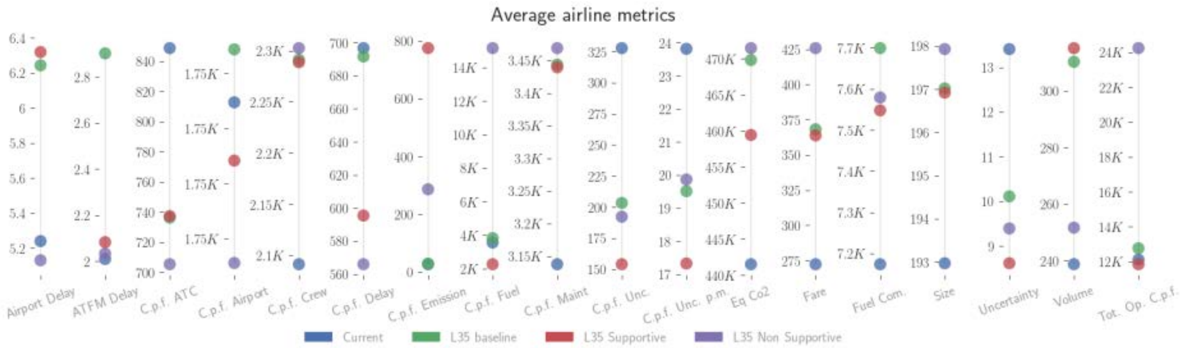

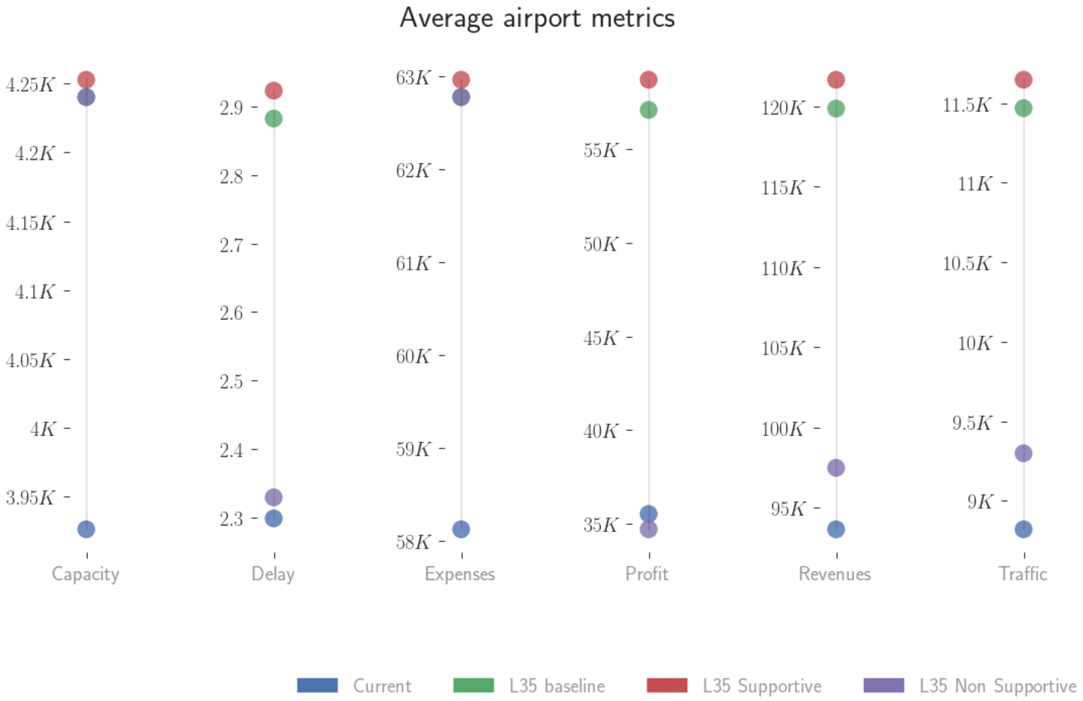

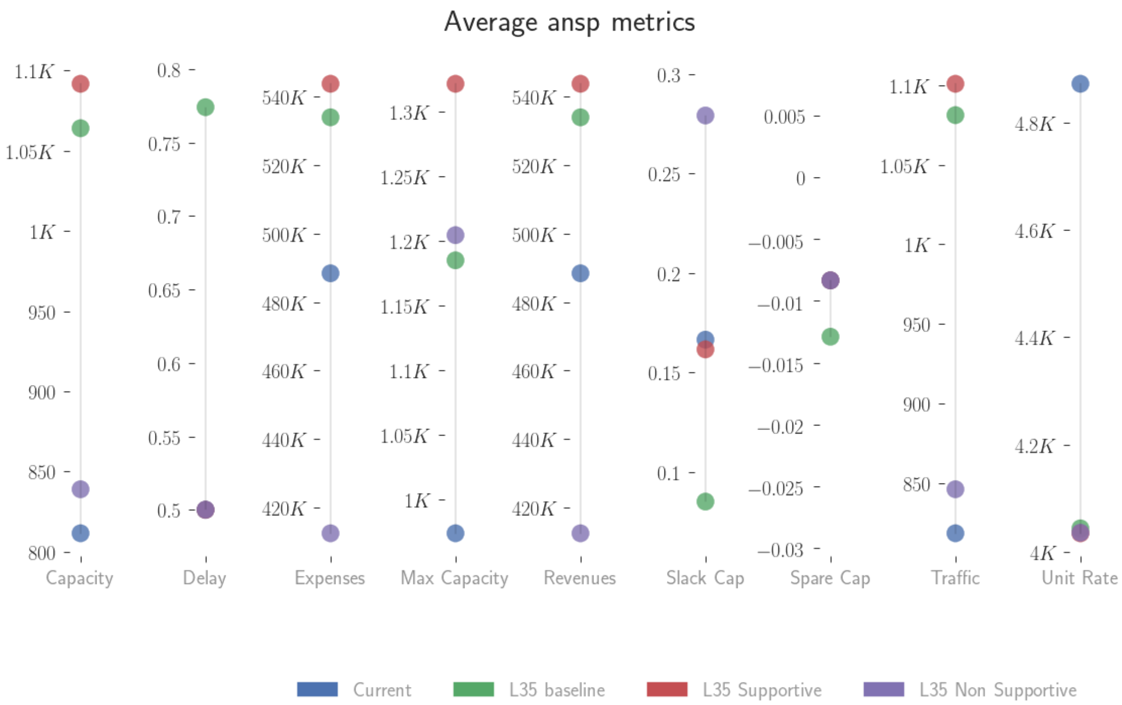
Average ansp metrics
Capacity
Delay
Expenses
Max Capacity
Revenues
Slack Cap
Spare Cap
Traffic
Unit Rate
Current
L35 baseline
L35 Supportive
L35 Non Supportive

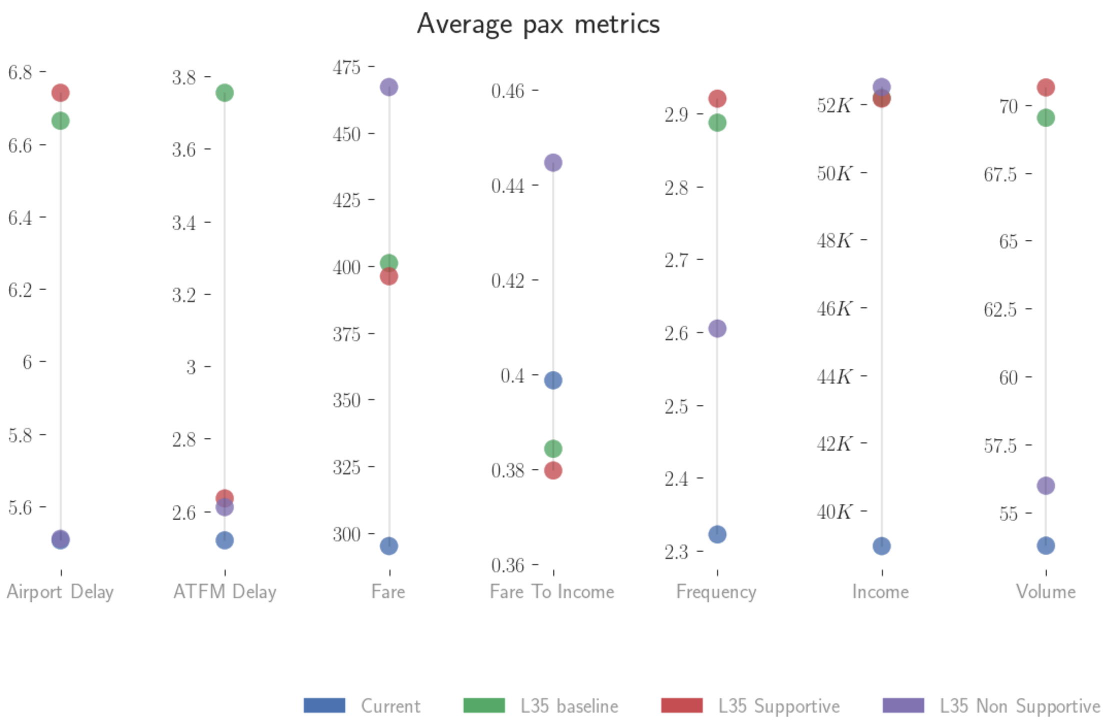
Average pax metrics
Airport Delay
ATFM Delay
Fare
Fare To Income
Frequency
Income
Volume
Current
L35 baseline
L35 Supportive
L35 Non Supportive

9.1.3.2.2 Supportive/Non-supportive in 2050

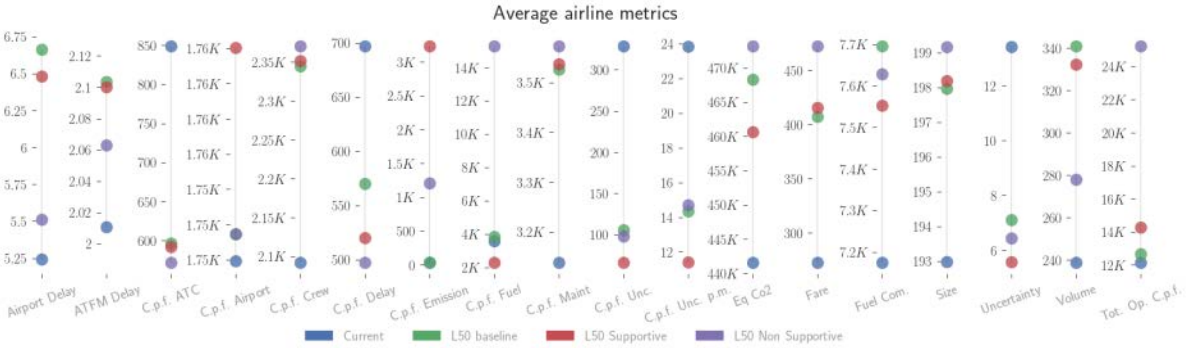

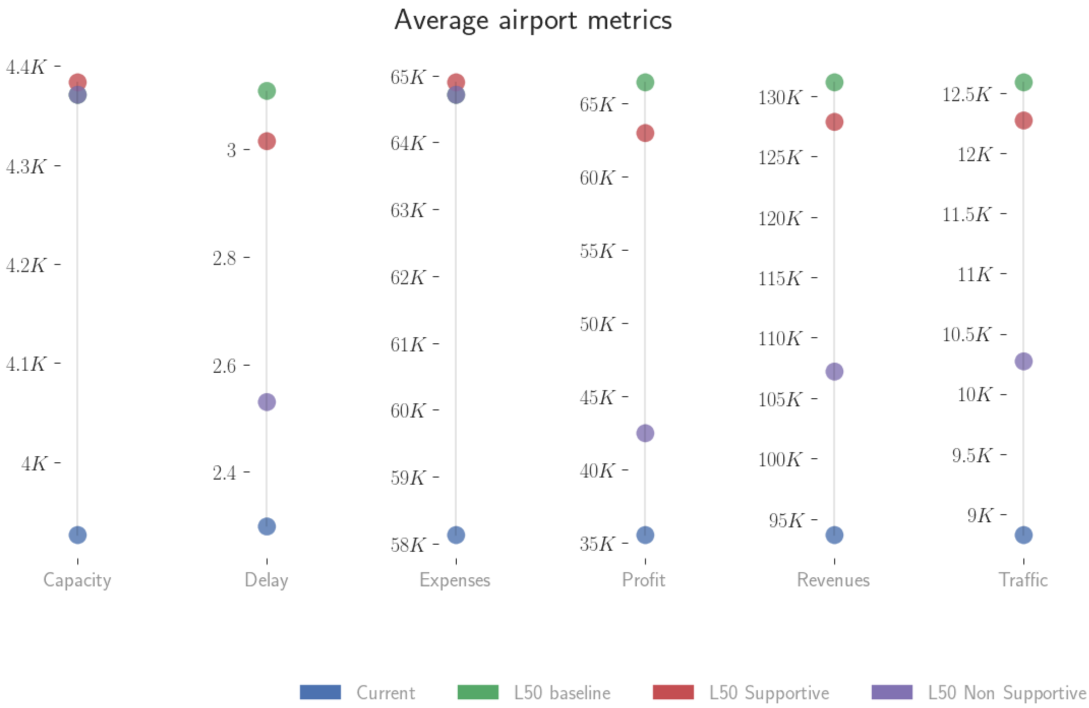

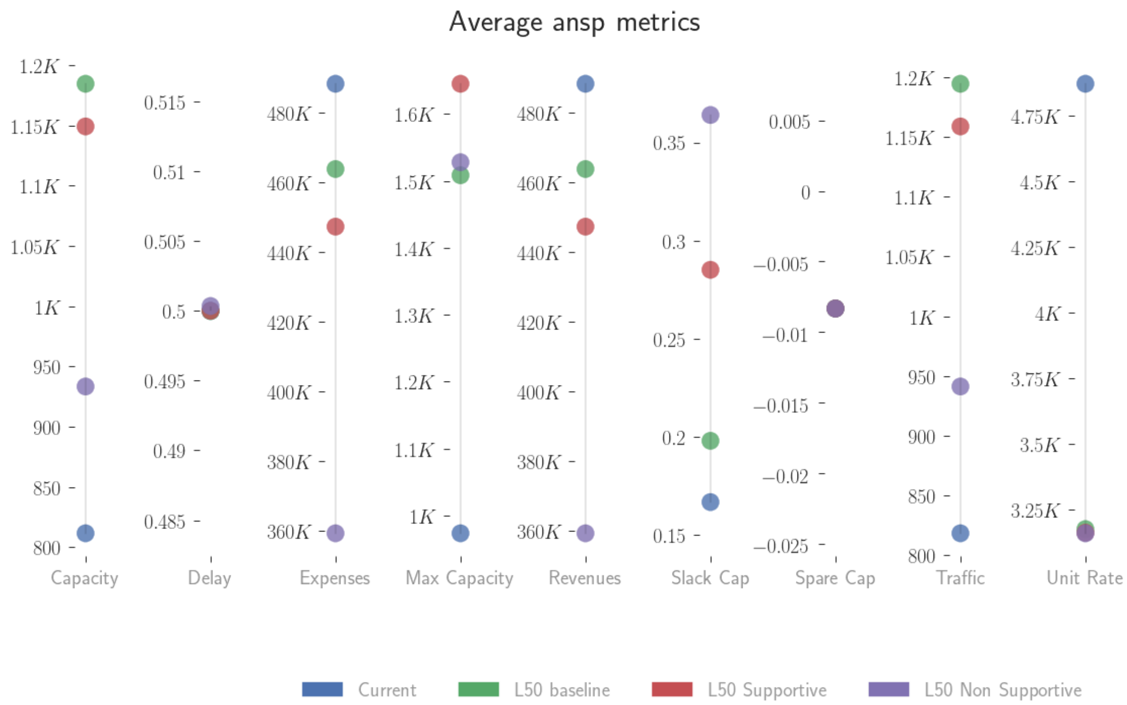

Average ansp metrics
Capacity
Delay
Expenses
Max Capacity
Revenues
Slack Cap
Spare Cap
Traffic
Unit Rate
Current
L50 baseline
L50 Supportive
L50 Non Supportive

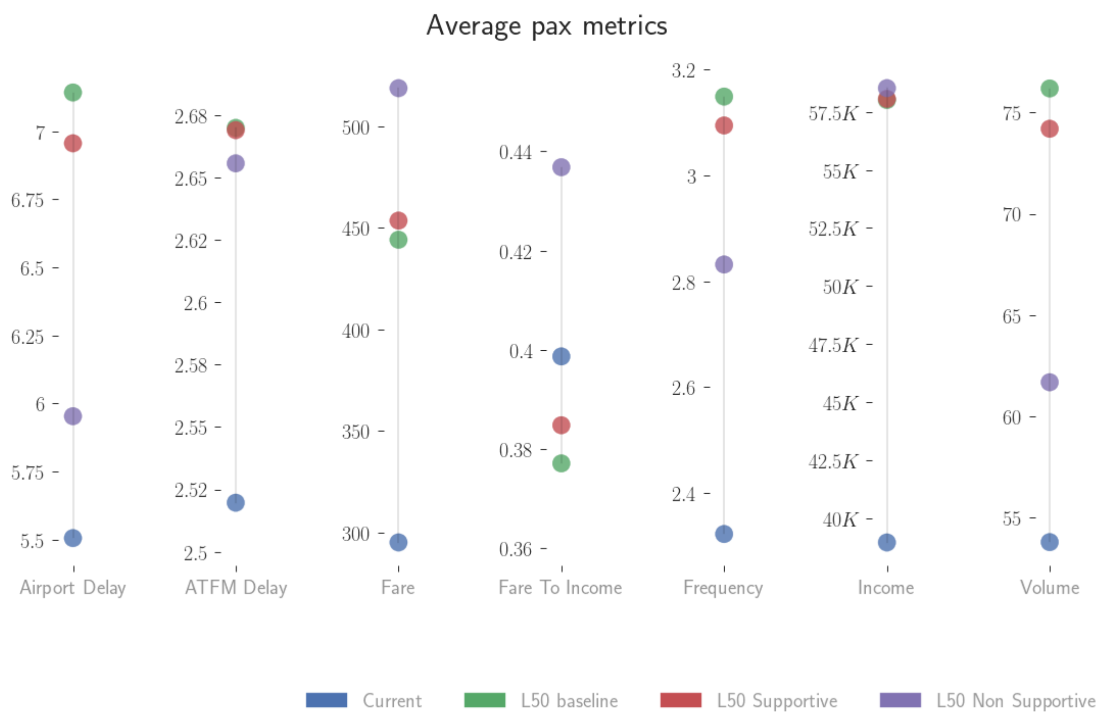

Average pax metrics
Airport Delay
ATFM Delay
Fare
Fare To Income
Frequency
Income
Volume
Current
L50 baseline
L50 Supportive
L50 Non Supportive

9.1.4 Cost structure

9.1.4.1 Baselines

9.1.4.1.1 Low baseline

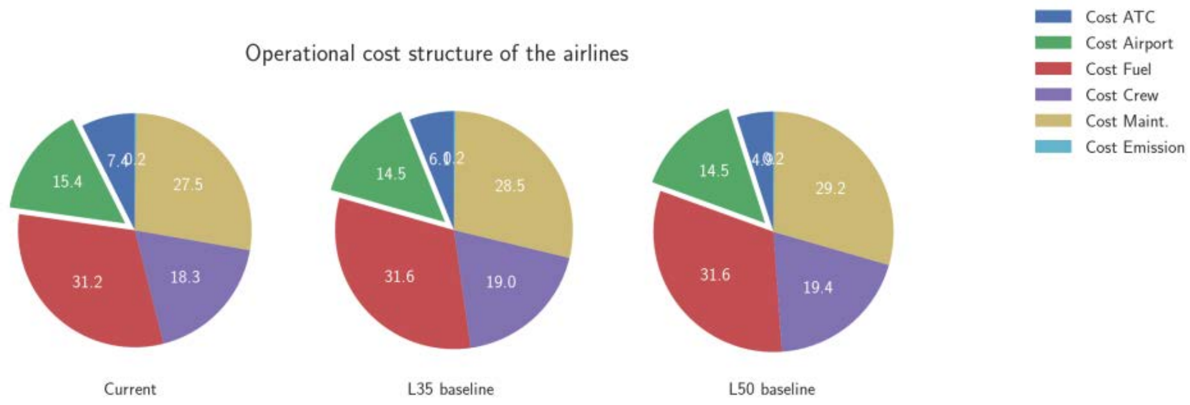

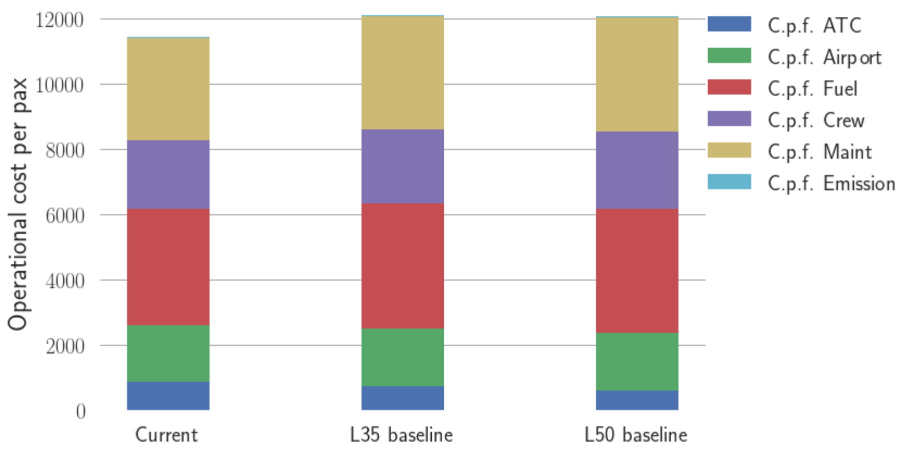

9.1.4.1.2 High baseline

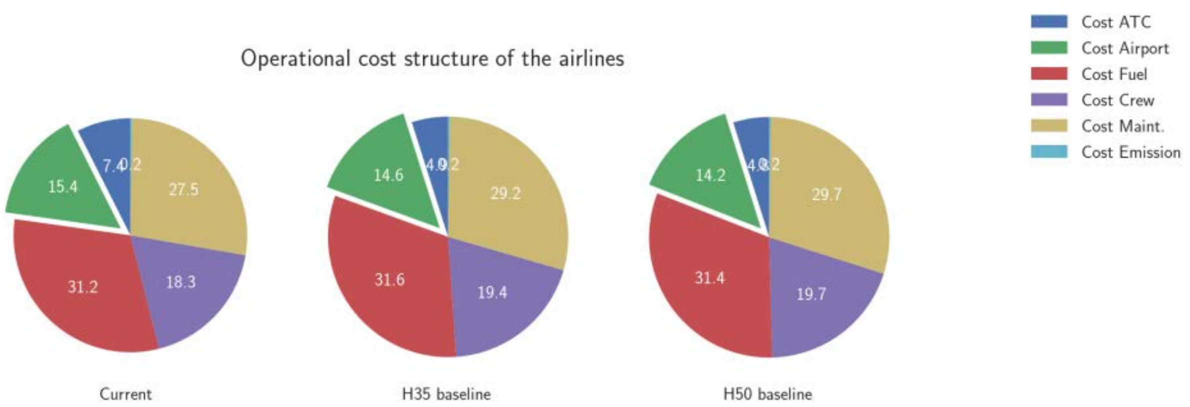

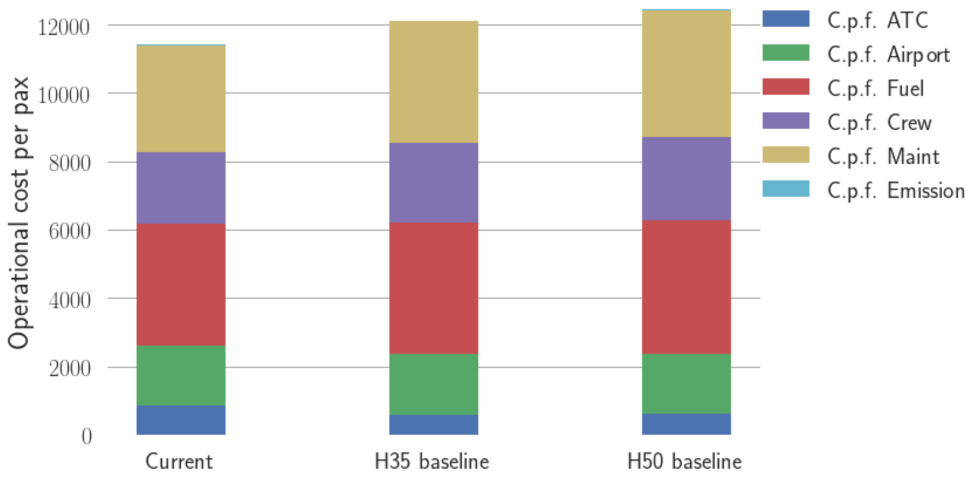

9.1.4.2 Supportive/non-supportive

9.1.4.2.1 Supportive/non-supportive in 2035

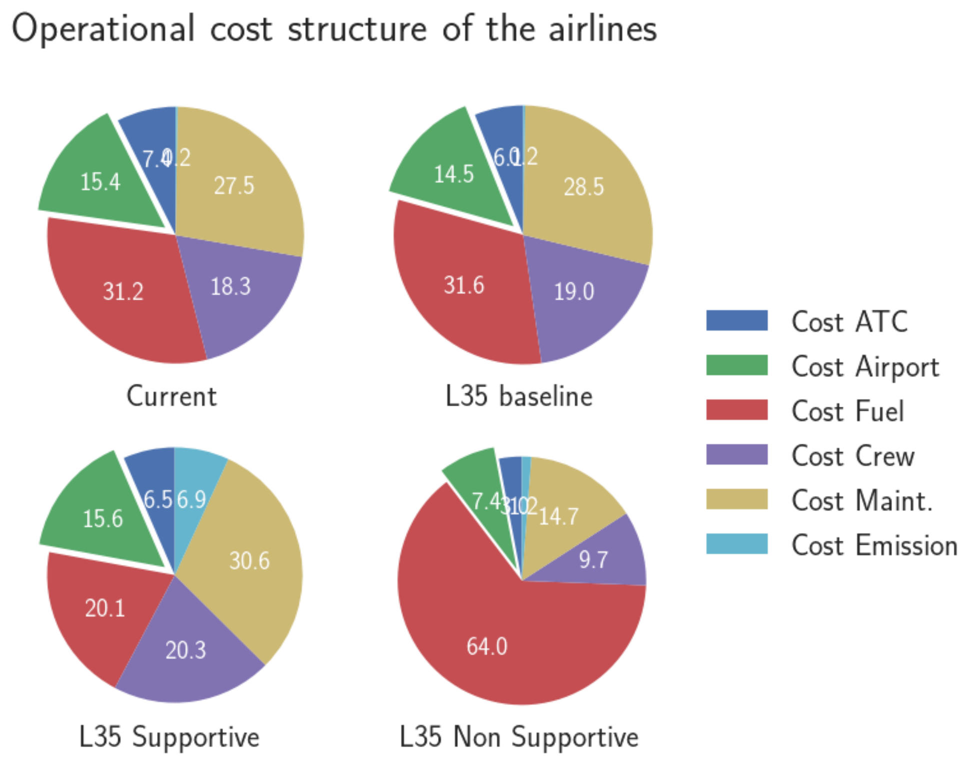

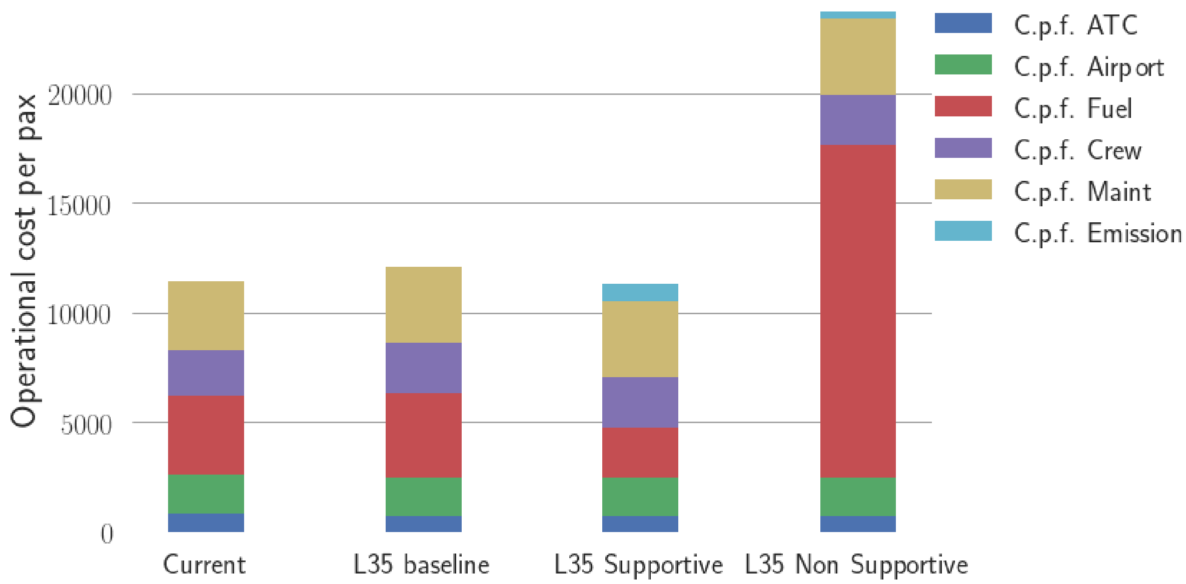

9.1.4.2.2 Supportive/non-supportive in 2050

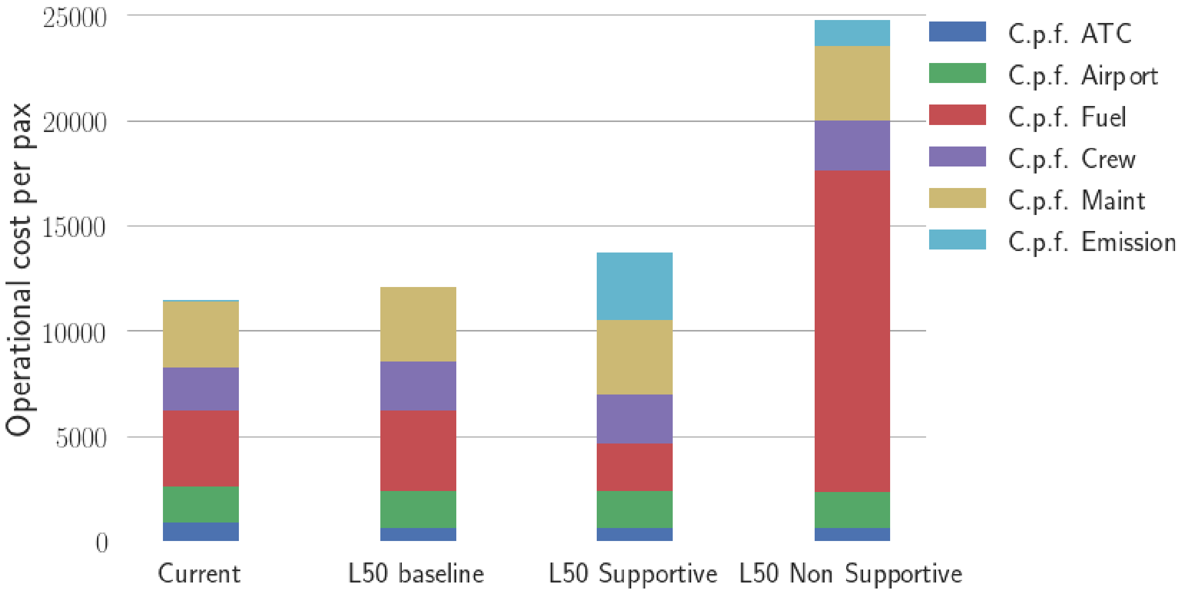

Operational cost structure of the airlines

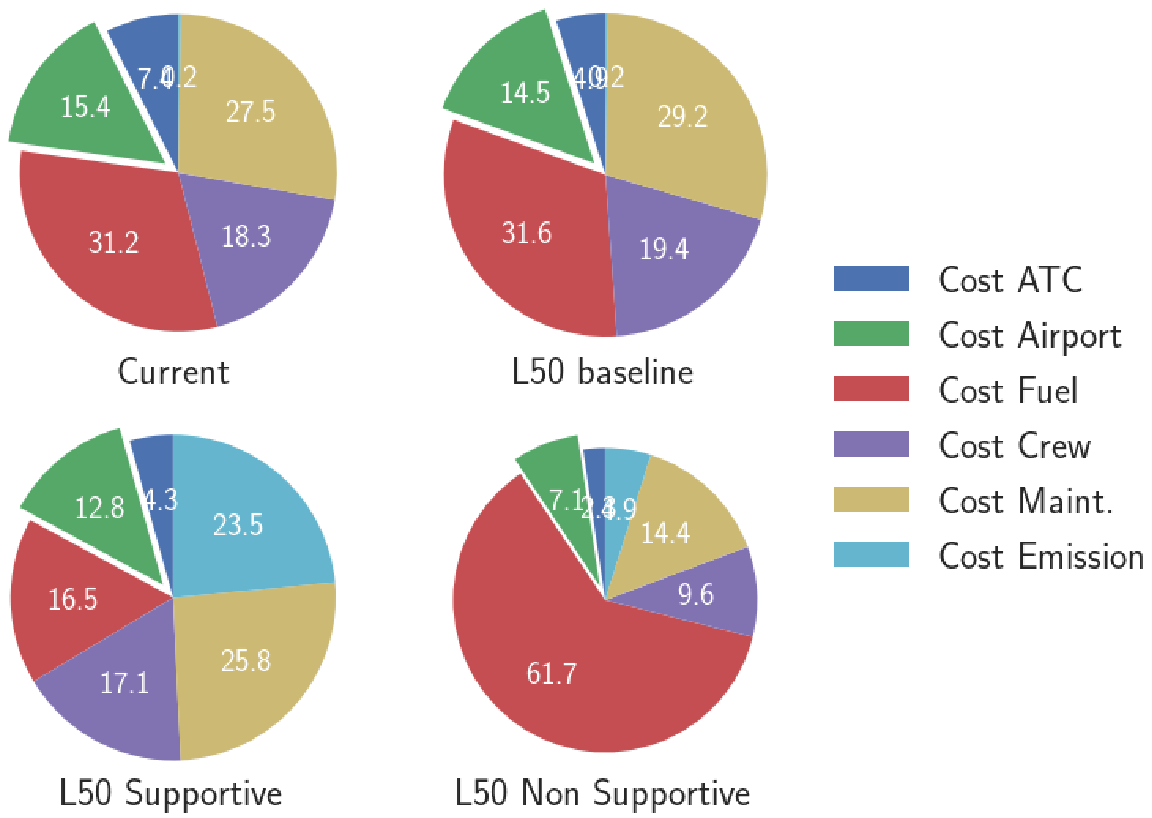

9.1.5 Statistical tests

9.1.5.1 Baselines

9.1.5.1.1 Low baseline

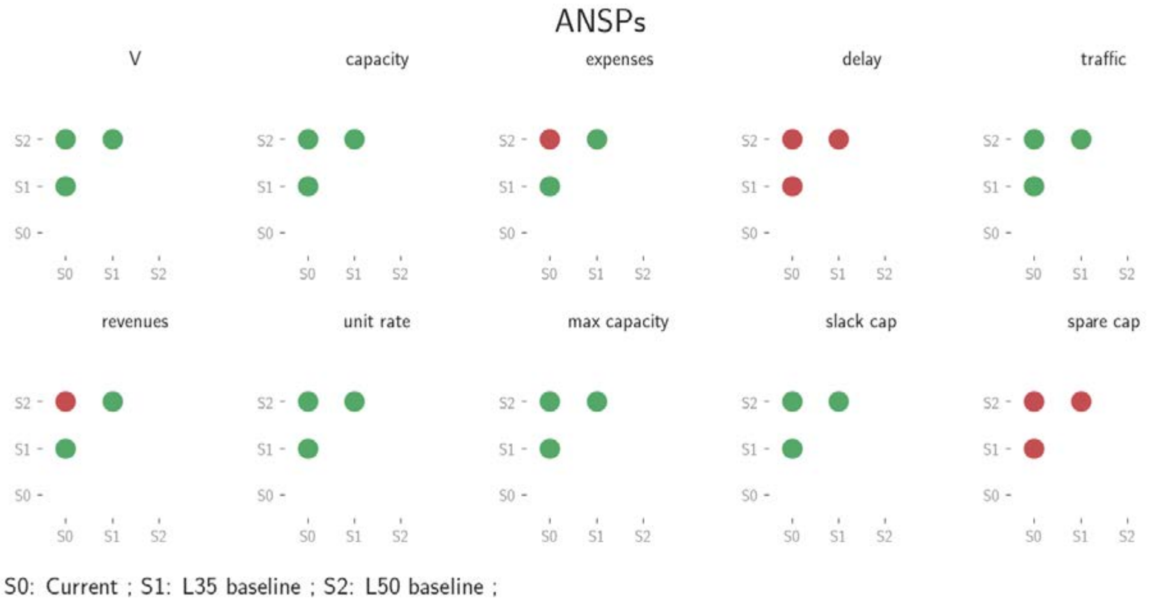

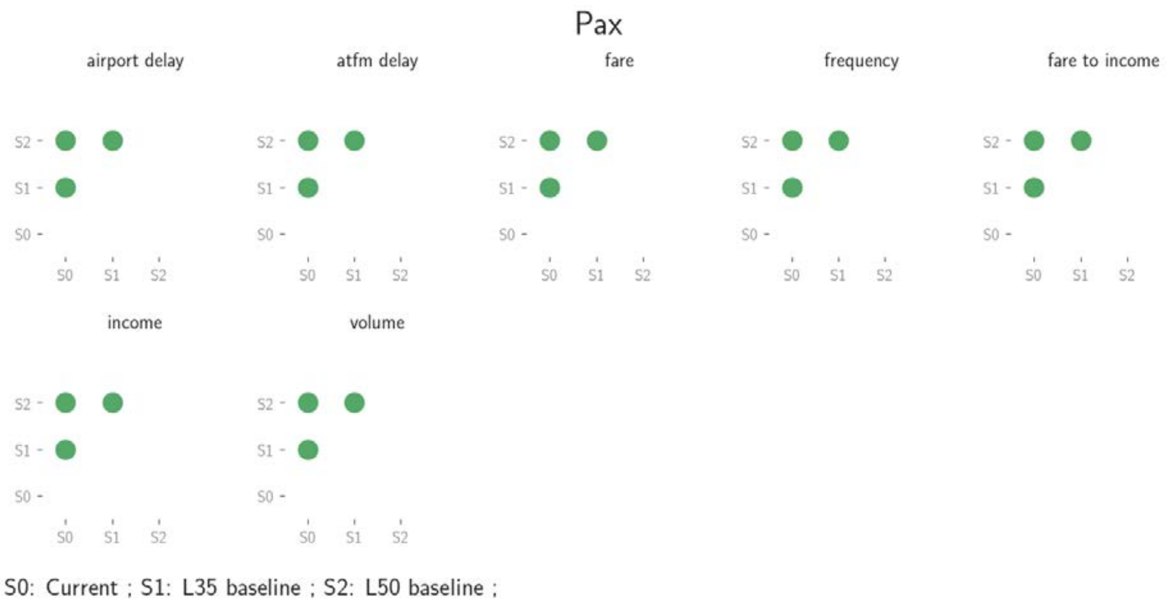

S0: Current ; S1: L35 baseline ; S2: L50 baseline ;

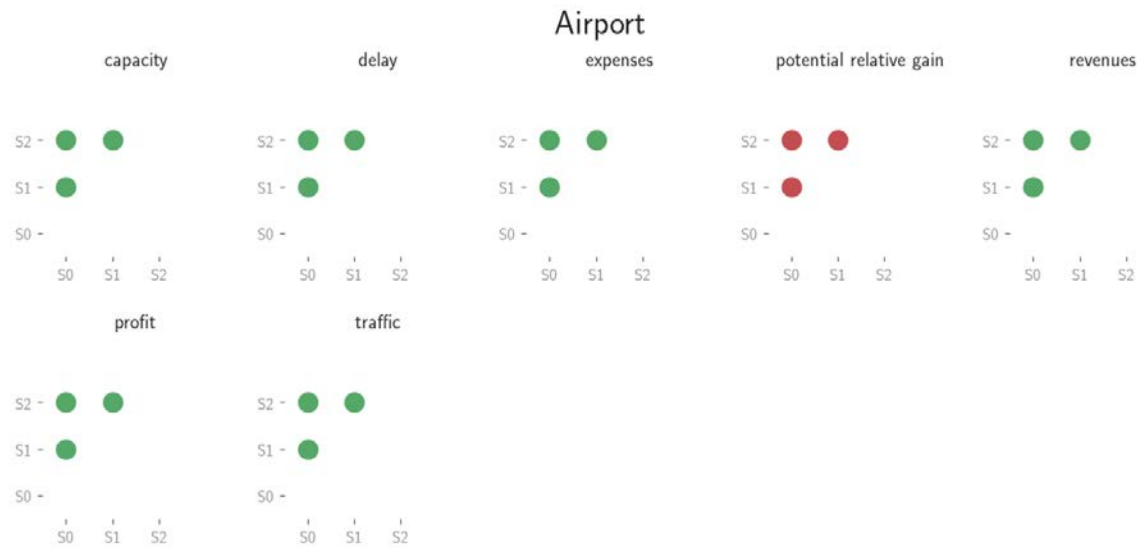

S0: Current ; S1: L35 baseline ; S2: L50 baseline ;

9.1.5.1.2 High baseline

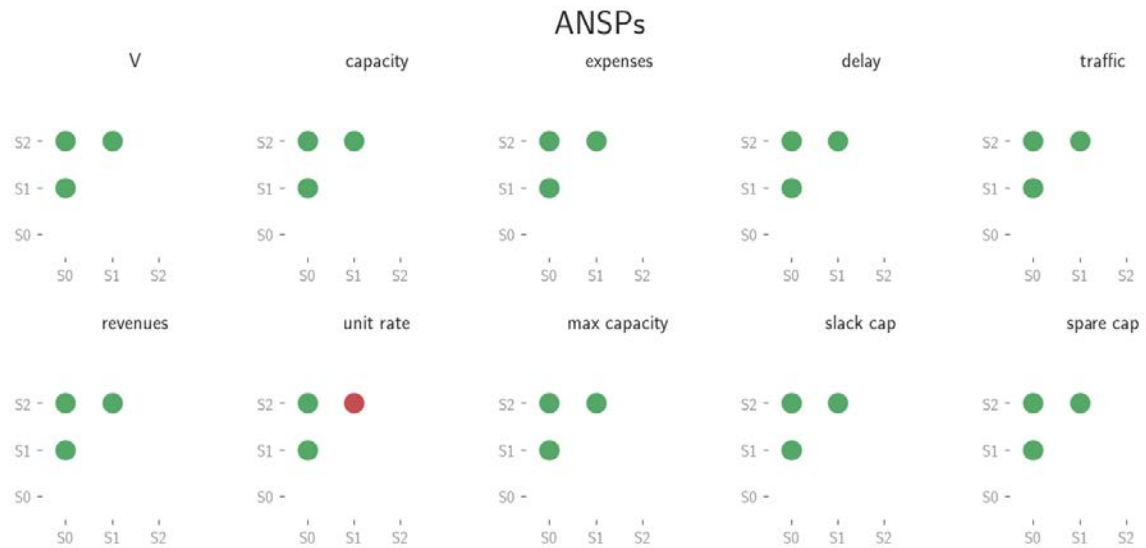

S0: Current ; S1: H35 baseline ; S2: H50 baseline ;

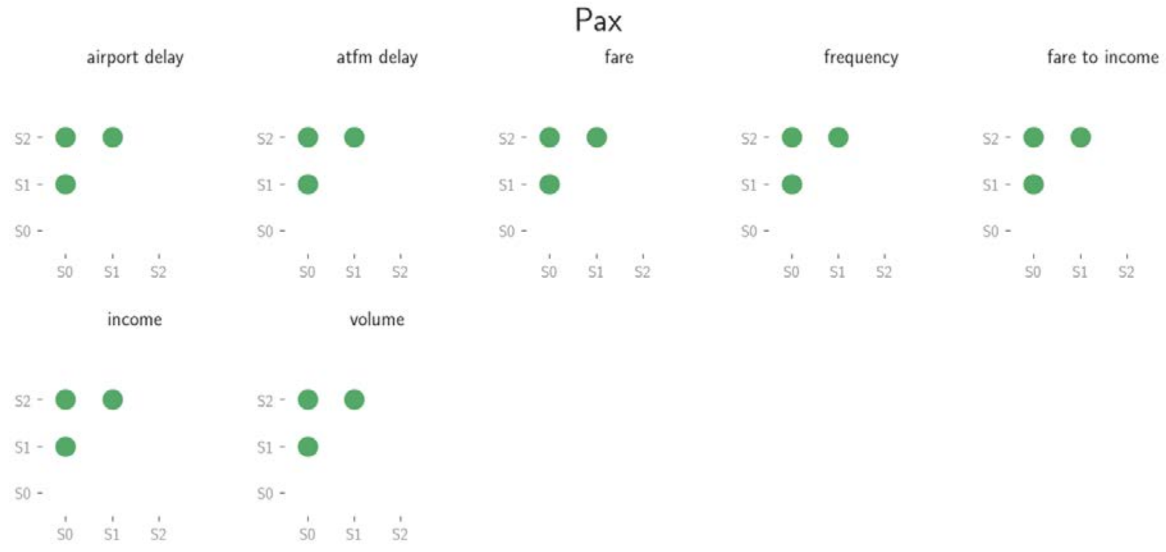

S0: Current ; S1: H35 baseline ; S2: H50 baseline ;

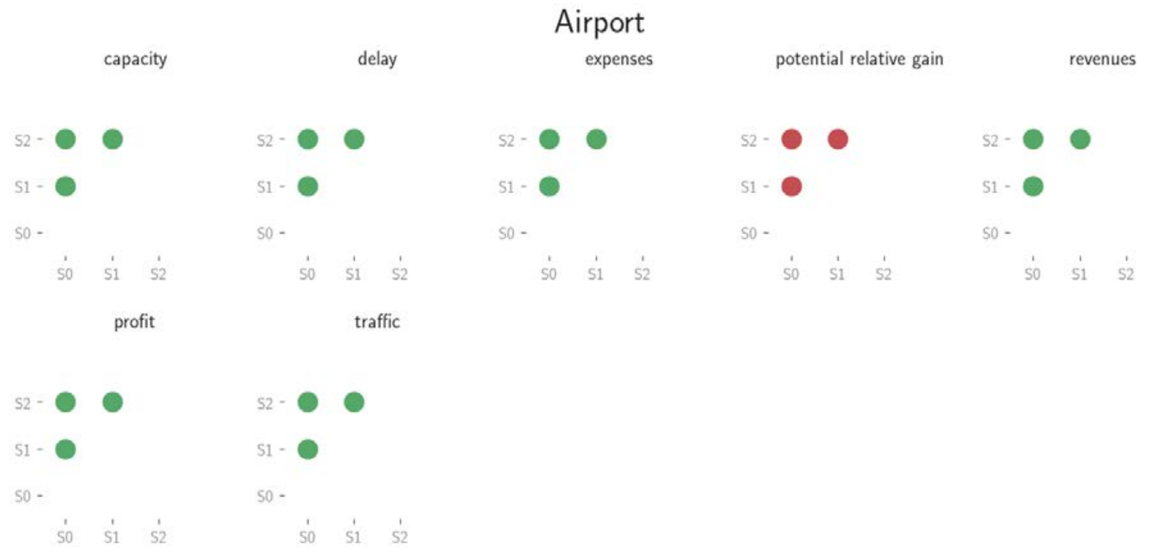

S0: Current ; S1: H35 baseline ; S2: H50 baseline ;

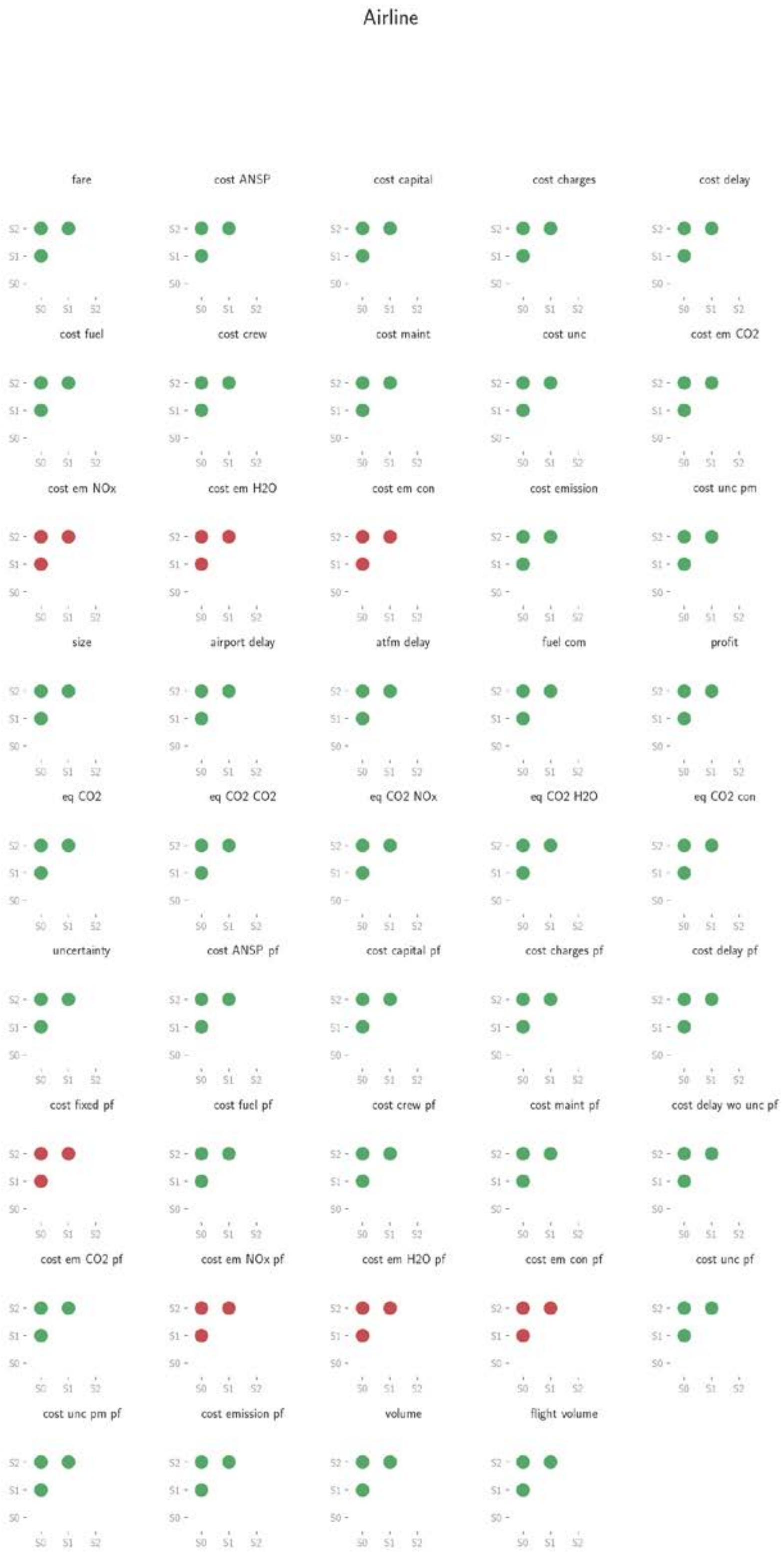

S0: Current ; S1: H35 baseline ; S2: H50 baseline ;

9.1.5.2 Supportive/non-supportive

9.1.5.2.1 Supportive/non-supportive in 2035

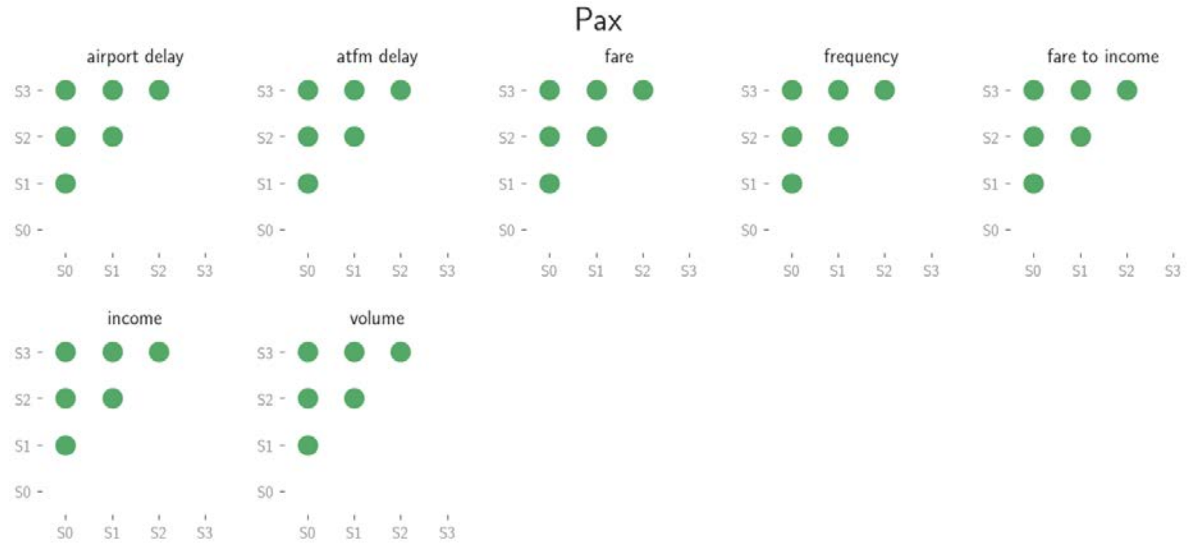

S0: Current ; S1: L35 baseline ; S2: L35 Supportive ; S3: L35 Non Supportive ;

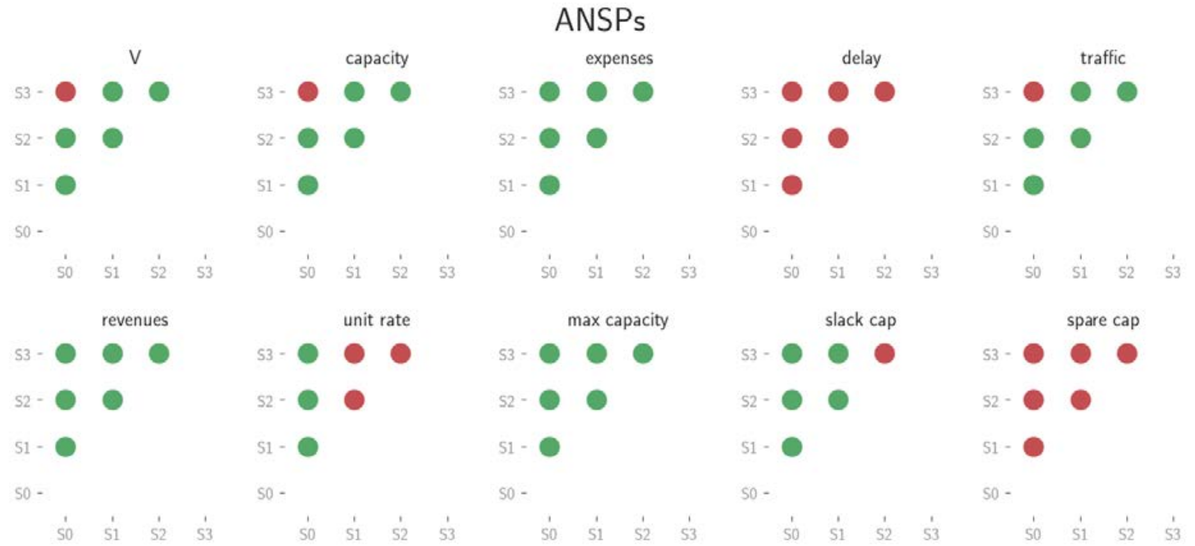

S0: Current ; S1: L35 baseline ; S2: L35 Supportive ; S3: L35 Non Supportive ;

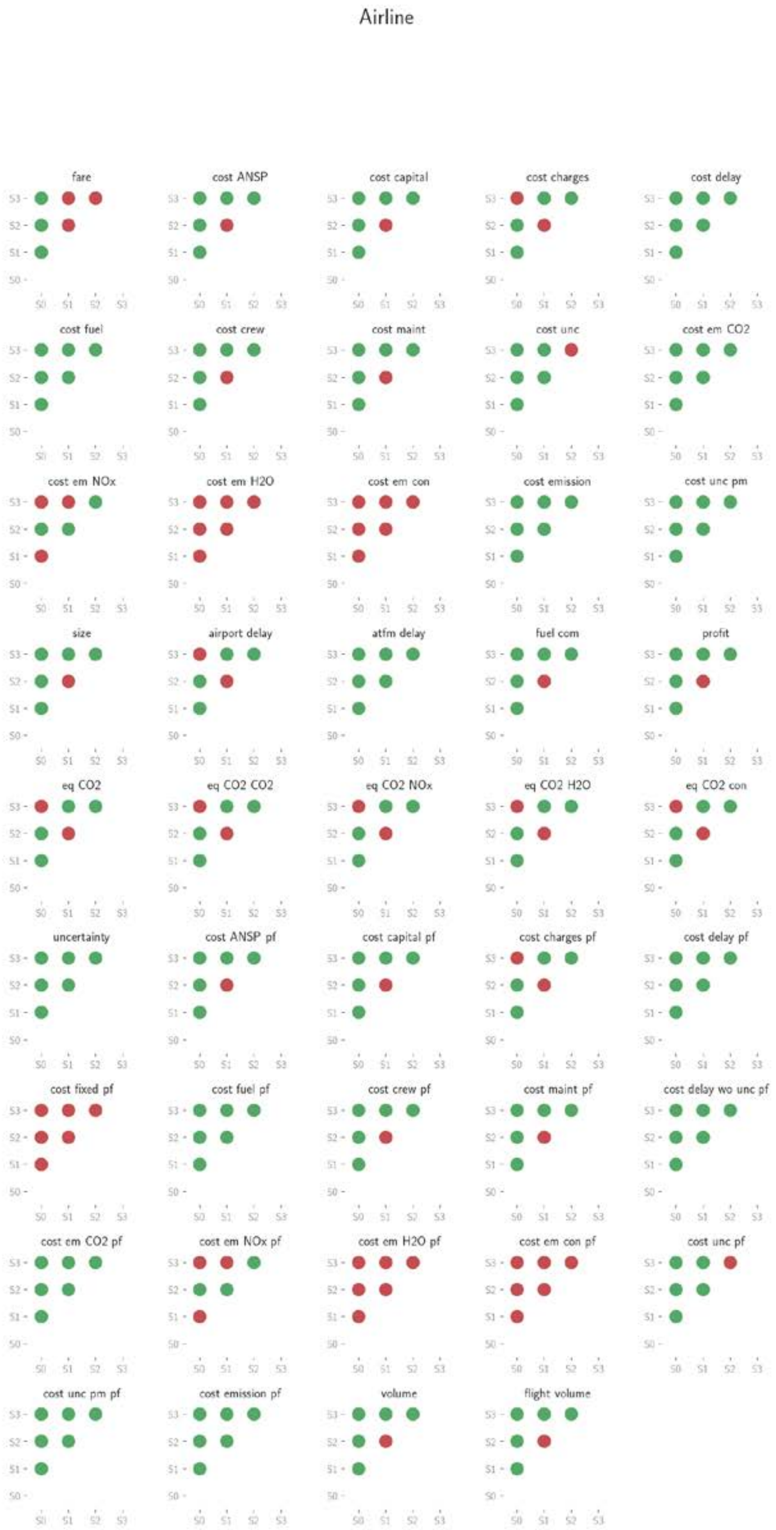

S0: Current ; S1: L35 baseline ; S2: L35 Supportive ; S3: L35 Non Supportive ;

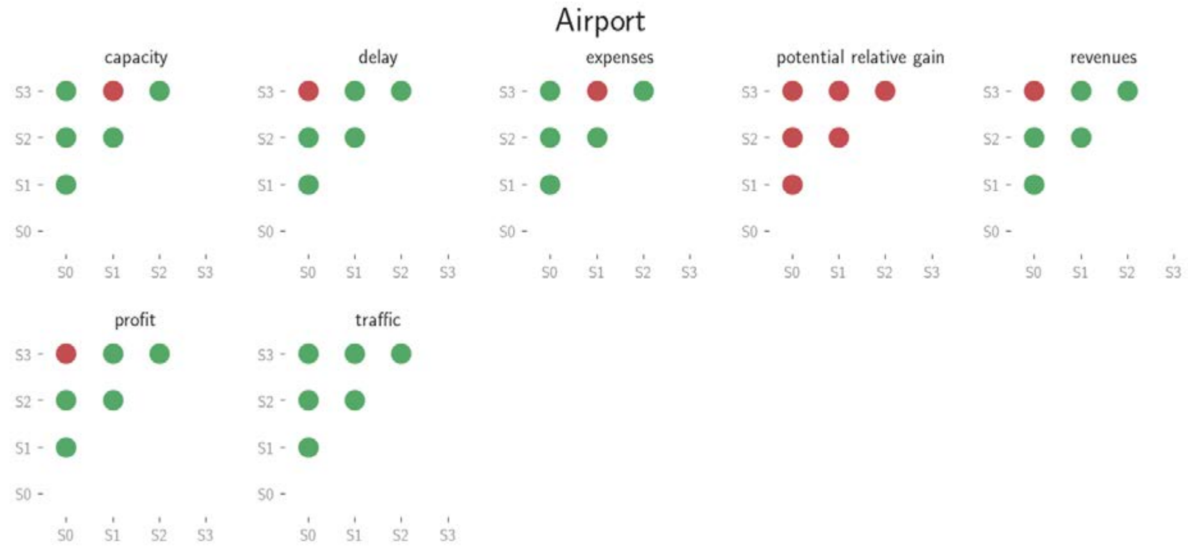

S0: Current ; S1: L35 baseline ; S2: L35 Supportive ; S3: L35 Non Supportive ;

9.1.5.2.2 Supportive/non-supportive in 2050

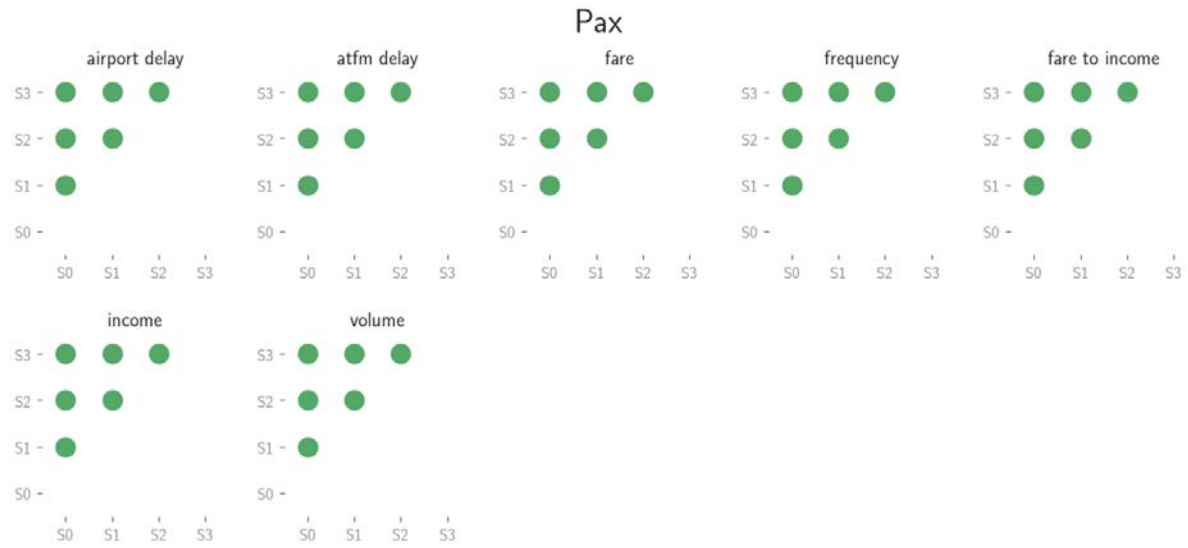

S0: Current ; S1: L50 baseline ; S2: L50 Supportive ; S3: L50 Non Supportive ;

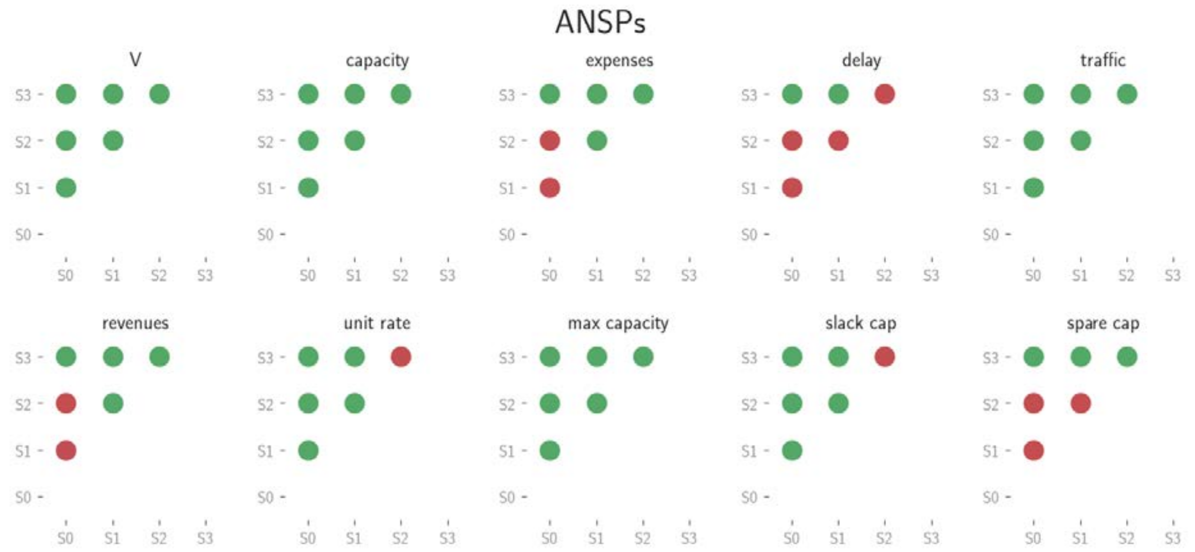

S0: Current ; S1: L50 baseline ; S2: L50 Supportive ; S3: L50 Non Supportive ;

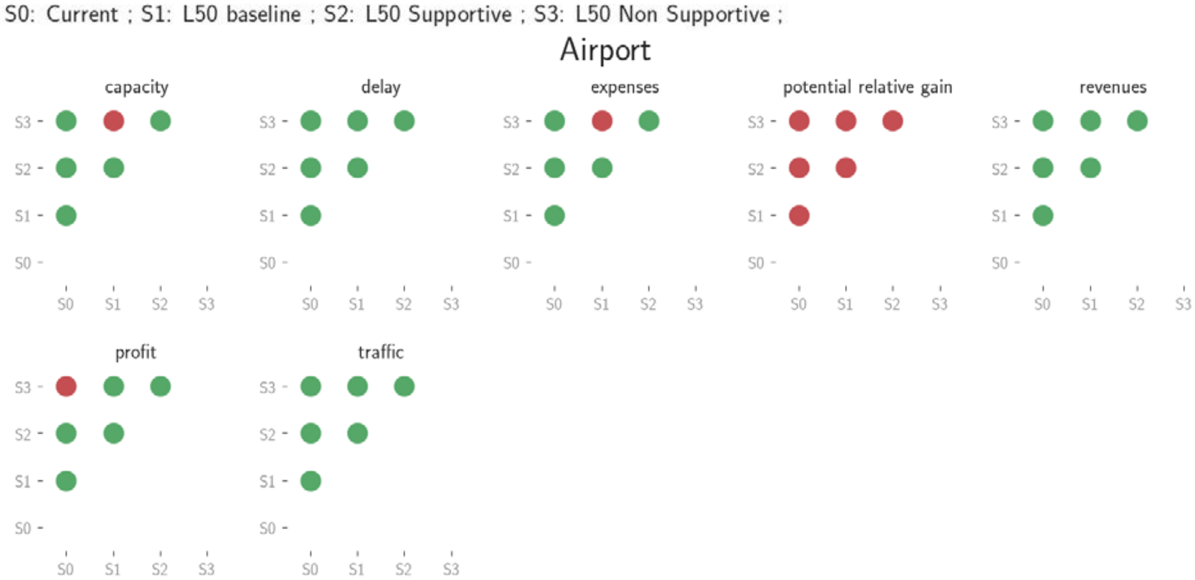

S0: Current ; S1: L50 baseline ; S2: L50 Supportive ; S3: L50 Non Supportive ;

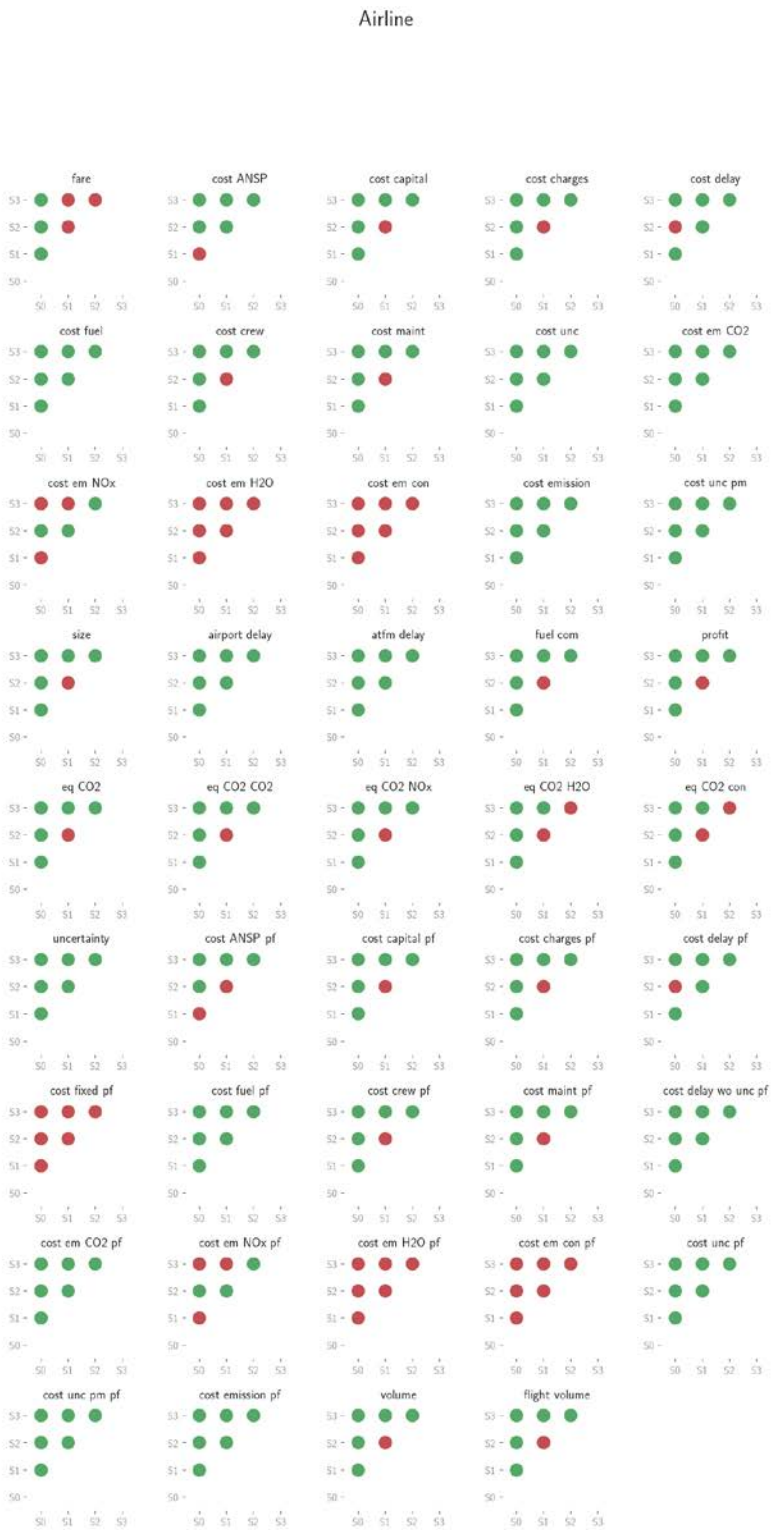

S0: Current ; S1: L50 baseline ; S2: L50 Supportive ; S3: L50 Non Supportive ;

10 Annex: Pre-tactical layer other capabilities results

This annex contains the results reported in D5.1 Initial assessment report for the pre-tactical layer to show further capabilities of this layer to do case studies analysis beyond generating the input to the tactical layer with the outcome of the strategic one.

In this case, two factors are considered in the pre-tactical layer. The factors considered affect the estimation of the cost of the different flight plan possibilities affecting the preferred route choice for the airlines. This will have an impact on the flight characteristics (distance, buffers, fuel usage) but also on the demand on the different ANSPs and on their revenue due to air navigation charges. The aim of selecting these factors is to show the model capabilities to reflect the impact of changes in the operating environment on the behaviour of the stakeholders and on the metrics.

10.1Factors modelled

10.1.1 Fuel prices (BEO1)

Fuel prices will have an impact on the total operating cost of the different flight plans leading to changes on the behaviour of the airspace users. If the value of fuel is low, the expected behaviour is that airlines will give higher preference to routes which might be longer but use cheaper airspaces, on the contrary, high fuel price will lead to more direct routes even at the cost of higher air navigation charges.

10.1.2 CRCO charges (BEO2)

The value of CRCO charges has been directly modified to analyse its impact on the route choice. Different options for the charges are considered: changing the size of the regions using the same rate, and the values of these rate. These changes might affect the incentive to select longer routes which incur in lower en-route cost. It is expected that as the size of the region which has unified rates increases, the airspace users will tend to select more direct routes reducing the total flight plan lengths and having an impact on the demand and revenues of the different ANSPs. For the calibration, the case when the value of the charges is set to zero is also modelled to see the impact of them on the route preferences. Another case considers an increase in double the cost of the charges which will incentivise the use of cheaper airspaces. Finally, a case where all the ANSPs which have currently a rate higher the average rate increase twice-fold while the airspaces that are below the average decrease it even further.

10.1.3 Factor and parameters summary

Table 56 shows the summary of the factors modelled in the pre-tactical layer with information on how the factor affects the different blocks and their expected impact in the model. In this case, all the factors are affecting the Flight Plan Generator block.

Table 56. Summary of factors modelled in pre-tactical layer

Factor	Summary description	Expected effect	Implementation of different values
BEO1 – Fuel prices	Fuel prices have an impact on cost of the different flight plan options for each flight.	Higher fuel values will lead to the preference of more direct and shorter routes even if the total CRCO cost is increased. Low price of fuel will have the opposite impact, encouraging airlines to select longer routes that use cheaper airspaces when possible. Note that the current value of fuel is already relatively low, for this reason it is expected that decreasing even further the cost of fuel might not have a high impact on the system.	LL: 0.1 EUR/kg L: 0.2 EUR/kg M:0.5 EUR/kg H: 1 EUR/kg HH: 2 EUR/kg

Factor	Summary description	Expected effect	Implementation of different values
BEO2 – CRCO charges	The CRCO costs of each flight route are computed based on the current rates. However, this factor will change the regions affected by similar rates.	When same CRCO charges are applied on a larger region there is less incentive to select longer routes as using airspaces with a lower rate might become harder. It is expected that as the region with unified charge increase there will be a reduction on the flight plan length of the flights selected.	Current: CRCO charges are computed as nowadays (See Section 2.2.1.3). Homogeneous rate: For all the charging zones which are charged by the CRCO an average unit rate is used. The only exceptions are the airspace of Lisbon Azores which is kept at is current rate and the communication charges of Shanwick airspace. This is maintained due to the characteristics of Oceanic traffic and airspace. The airspaces that are not part of the CRCO charging scheme are maintained as in the current scenario. FABs: In this case an average unit rate is computed per FAB. Moreover, the charges are not computed per NAS but at a FAB level, i.e., considering the entry and exit point within the FAB and not within the airspaces of the countries that are part of the FABs. Once again, note that Lisbon Azores, Spain Canary Islands and Shanwick airspace are considered outside the FABs. The airspaces that are not part of a FAB keep their rate as in the current scenario. Zero: In this case EUROCONTROL's CRCO's member states rates are set to zero, in order to better understand the impact of en-route airspace charges to the route preference. Double: All rates of EUROCONTROL's CRCO's member states rates are double. This increase the relevance of en-route airspace costs. Imbalance: In this case, airspaces with a unit current rate that is higher than the average double their rate while countries with a unit rate lower than the average decrease it by half. This generates higher incentives to use cheaper airspaces.

10.2 Scenarios and results

The different scenarios tested by the pre-tactical layer focus on how changes on fuel and en-route airspace charges have an impact on the flight plans that are generated. The changes of those factors impact the operating cost of the different flight plans and routes options and hence it is expected to have an impact on the route selection preference by airlines. This in its turn will impact the characteristics of the routes (e.g., the length of the flight plans or the buffer time) and the demand and revenue of the different NAS.

The metrics that are quantified at the pre-tactical layer include:

- For flight plans:
    - Distance (NM);
    - Time (minutes);
    - En-route airspace cost (euros);
    - Estimated flight plan fuel usage (kg);
    - Buffer and taxi times (i.e., block time - flight plan time) (minutes).
- For NAS:
    - Revenue due to airspace en-route charges (EUR);
    - Demand (aircraft entry count).

The Route Generator and Trajectory Generator have been executed once, the Flight Plan Generator has been run once per scenario and the Flight Plan Selector has been executed 50 times per scenario.

The results obtained have been analysed and presented below as:

- Impact of fuel cost on flight plan selection;
- Impact of en-route airspace charges on flight plan selection;
- Impact of fuel cost on revenue per NAS;
- Individual origin destination examples of route selection and impact on NAS;
- Demand over time at NAS.

10.2.1 Impact of fuel cost on flight plan selection

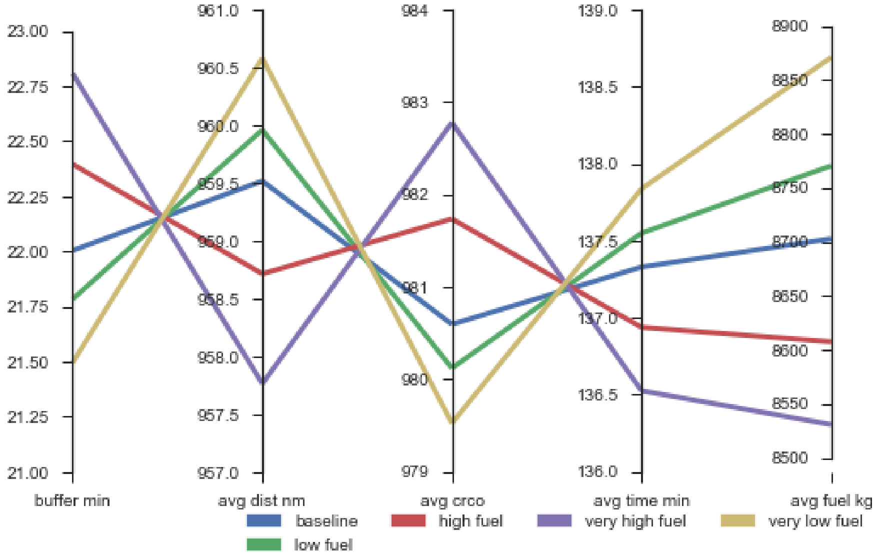

Figure 115. Impact of fuel cost on FP selection characteristics

Figure 115 presents the impact of fuel cost on the different metrics computed for the flight plan characteristics. As previously explained, the Flight Plan Selector has been executed 50 times per scenario, hence the values shown are the average per flight of all the values obtained for all the flights in the Current scenario (based on the schedules of 12SEP14).

For all the metrics, the changes in their values for the different scenario are as expected. As fuel cost gets higher, the flight plans selected tends to be on average shorter to save fuel even if incurring on higher en-route airspace costs. This leads to, in average, shortest flight times which in its turn increases the time scheduled for taxi and as buffer (buffers in the figure).

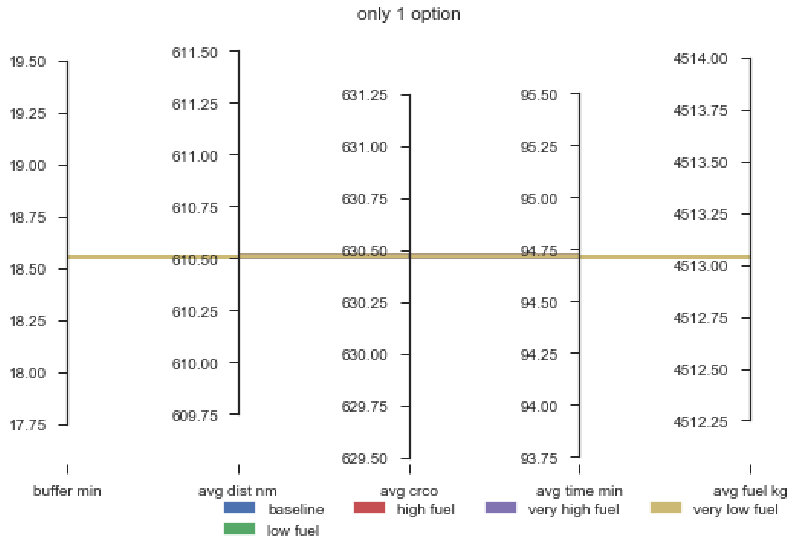

Figure 116. Impact of fuel cost on FP selection characteristics; flights with one route option

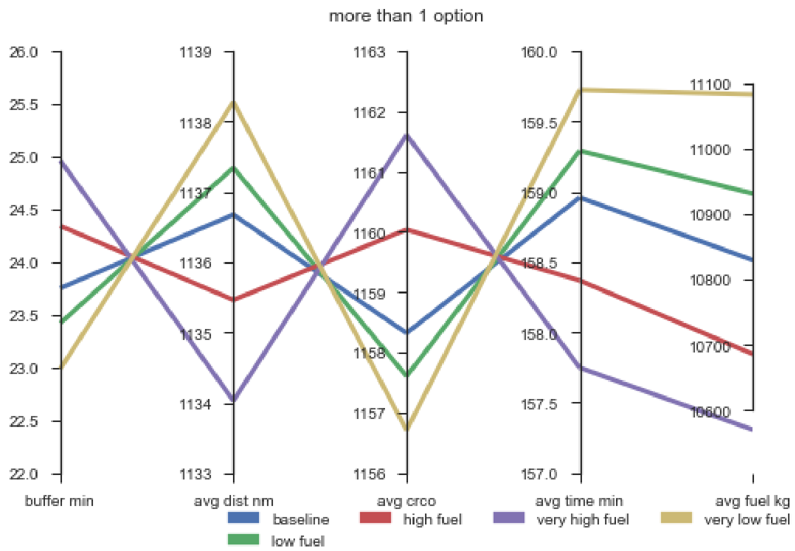

Figure 117. Impact of fuel cost on FP selection characteristics; flights with >1 route option

If for a given flight plan there is only one possible route, the changes on fuel cost will not modify the results (as shown in Figure 116). The Flight Plan Selector does not have any more options to choose from. On the contrary, if only flights with more than one option are considered, the differences in the metrics for the different scenarios are even larger (see Figure 117).

10.2.2 Impact of en-route airspace charges on flight plan selection

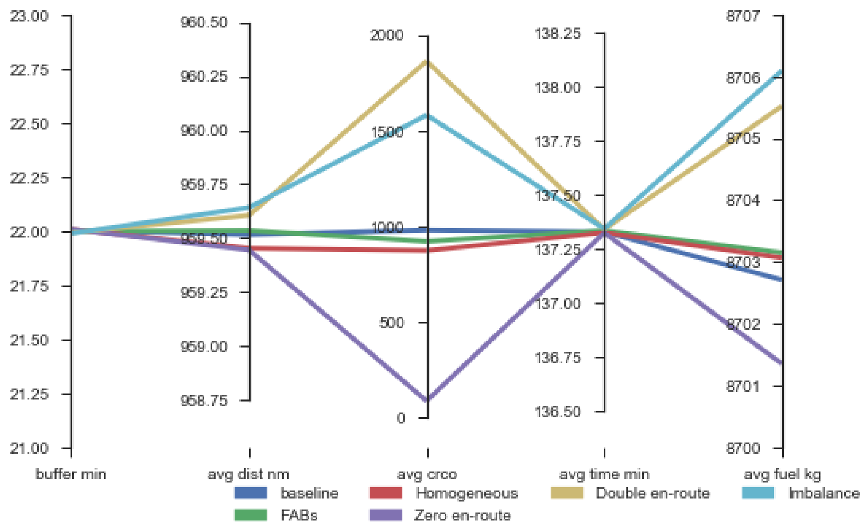

Figure 118. Impact of en-route charges on flight plan selection characteristics

Figure 118 presents the effect of modifying the airspace en-route charges on the different flight plan metrics. The results obtained, so far, show that at an average level the impact of modifying the regions that are used to charge or even the en-route airspace unit rate is very small. Their impact is more limited that changes on fuel price.

This might be due to the fact that the presented results are averaged for all the flights in the ECAC region. In order for changes in fuel having an impact on the flight plan selection due to fuel, only having more than one flight plan between a given origin and destination is required. However, for changes to the en-route charges, those different fight plans must also use different airspaces with different unit rates, and those price differences be high enough to drive the flight plan selection. This is not the case for the majority of flights, as for example, all fights inside a given NAS or between adjacent NAS, usually only have routes which use the same airspace. Therefore, the effect of this factor is not evenly distributed across the ECAC region and focus should be given to specific origin and destination pairs where different options using different en-route airspace charges regions are possible. This will be presented in Section 10.2.4.

10.2.3 Impact of fuel cost on revenue per ANSP

A side effect of changing the cost of fuel is that, as airspace users change their preferred trajectories, the amount of service units provided by each airspace also change. This has a direct impact on the demand over the different regions and also on the revenues obtained from en-route airspace charges.

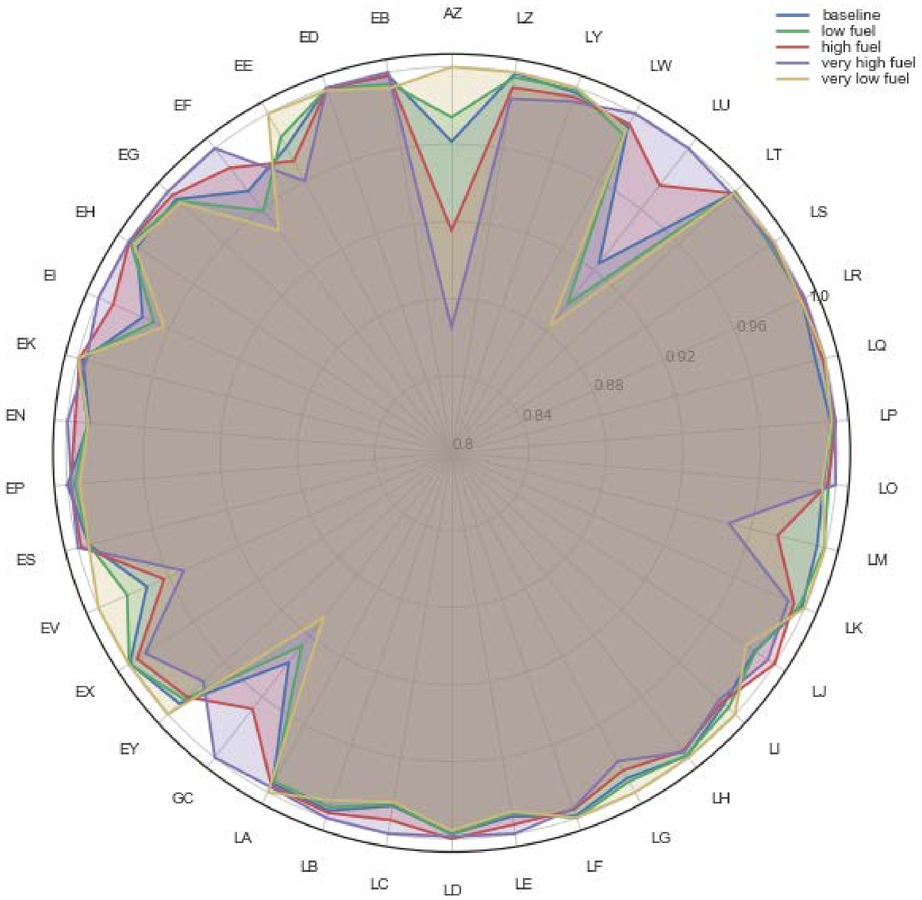

Figure 119. En-route airspace usage revenues per NAS for different fuel scenarios

Note AZ: Portugal Azores airspace, EX: Shanwick airspace

Figure 119 shows for each airspace charging zone the change in revenues obtained with respect to the maximum revenue observed in that charging zone. In some cases, changes in fuel cost might represent significant changes on total revenues obtained. For example, AZ region (Lisbon Azores airspace) receive the maximum revenue due to en-route charges in the scenario where fuel price is very low and as price of fuel increases the revenue decreases, dropping to 87% with respect to the revenue obtained in the very low fuel scenario when fuel price is set to very high. This is consistent with more flights selecting shorter, more direct routes to-from the Canary Island when cost of fuel is higher, but more flights using Portugal Oceanic Azores airspace when fuel is cheaper to save on the airspace en-route service cost.

10.2.4 Individual origin destination examples

As shown in the previous sections, if metrics are computed and averaged at an ECAC level, the impact of some of the factors might not be appreciated. Changes in fuel and, particularly, changes in airspace en-route service charges have special effect on specific origin and destination areas where airspace users have a larger selection of possible routes and with routes that cross different charging zones which have different unit rate. Considering this, in this section, the effect of the modelled factors on a selection of origin and destination pairs are presented in more detail.

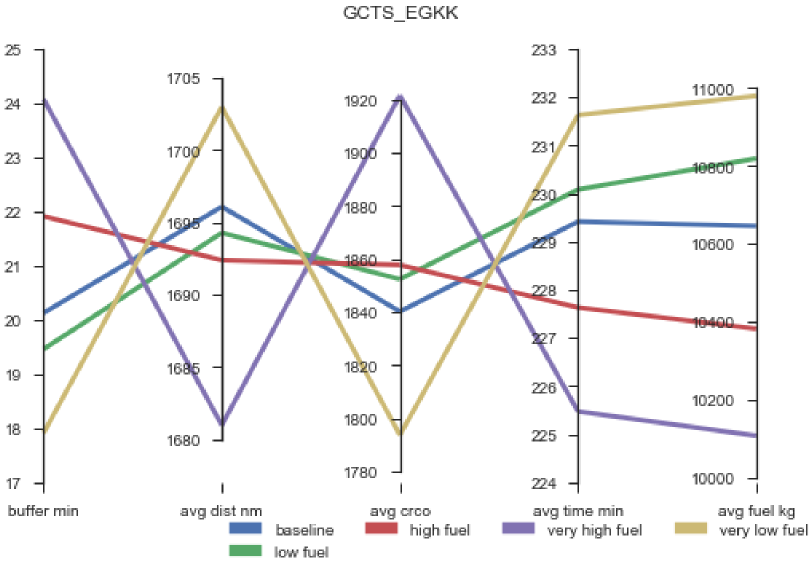

Figure 120. Impact of fuel cost on FP selection characteristics for GCTS to EGKK flights

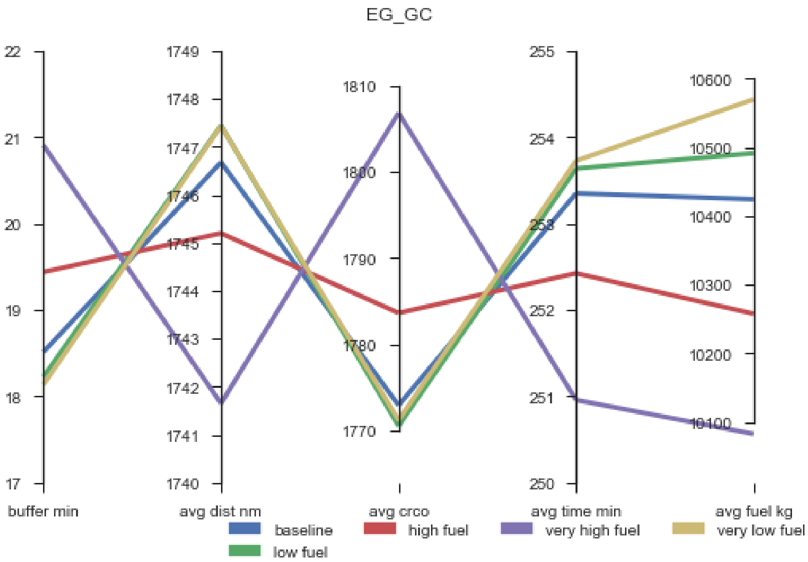

Figure 121. Impact of fuel cost on FP selection characteristics for all EG–GC flights

In flights between the United Kingdom (EG) and the Canary Island (GC), sometimes, there is the possibility of trading longer routes and using cheaper airspaces by using Portugal Azores and Oceanic Shanwick airspace (see Figure 126.a to see all possible route between those two regions).

Figure 120 and Figure 121 show the impact of fuel on the flight plan selected for the flights between GCTS and EGKK and for all the flights between EG and GC respectively. Note that in general, as the current cost of fuel used is relatively low (0.5 EUR/kg), decreasing its value even further has a smaller impact on the route selection than increasing it. The metrics observed in the high fuel scenario diverge from the values observed for the other scenarios as expected: with higher fuel costs shorter, more direct routes are preferred even at the cost of higher airspace en-route charges. This has an impact on the time allocated for taxi and buffer which, in average, gets increased by 3 minutes.

In Section 10.2.2, where the analysis of the changes to airspace charges was presented, it was stated that the impact of changes on airspace charges was small on the flight preferences. However, that is the case when all flights are considered, when looking at the results for only the flights between GCTS and EGKK and for those between EG and GC the impact of this factors is more observable, even if still small.

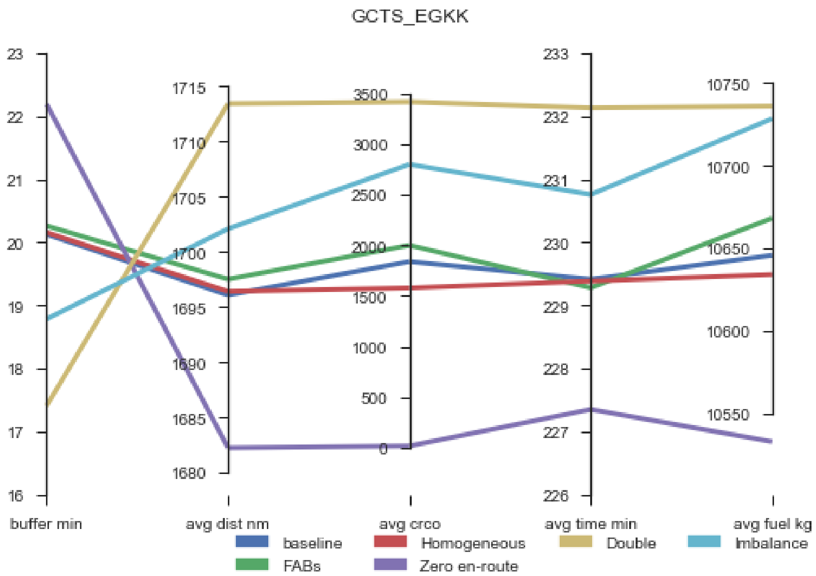

Figure 122. Impact of en-route charges on FP selection for GCTS to EGKK flights

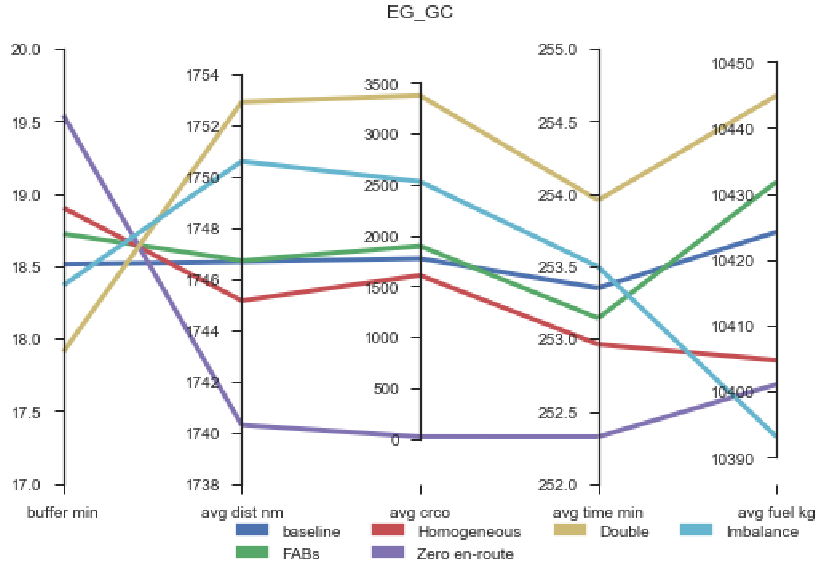

Figure 123. Impact of en-route charges on FP selection for all EG–GC flights

When en-route airspace charges are increased or when there is a larger imbalance between regions, longer cheaper routes are preferred; whilst when the cost of the airspace is removed (scenario Zero en-route) then shortest direct routes are submitted for the flight plans (see Figure 122 and Figure 123).

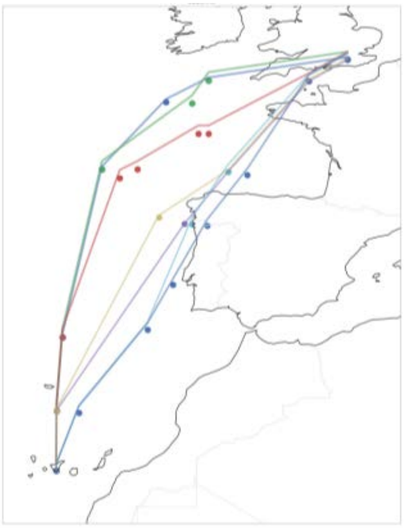

a) Baseline scenario

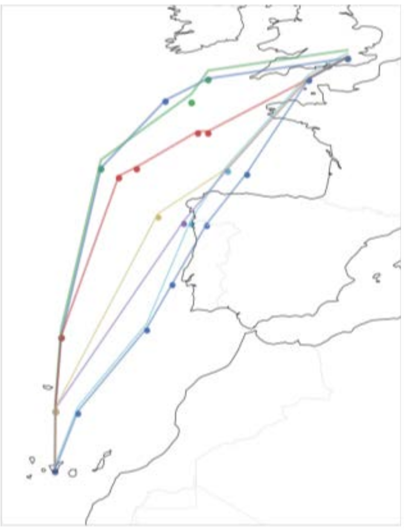

b) Very low fuel

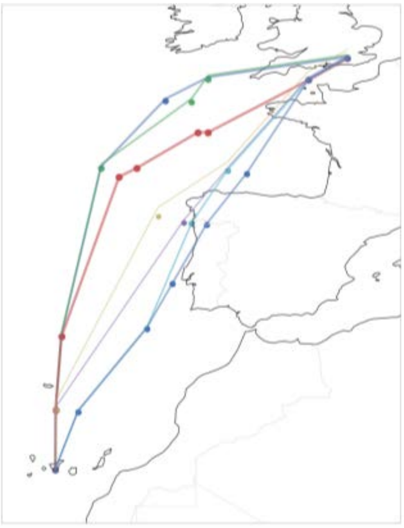

c) Low fuel

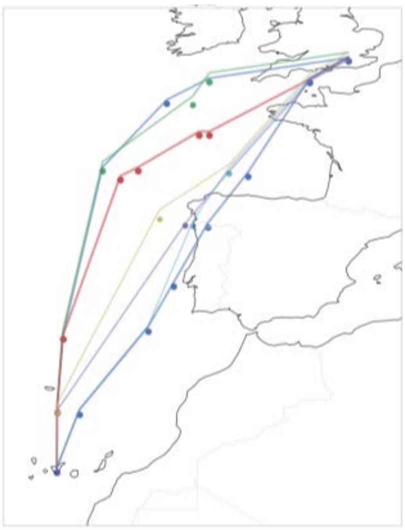

d) High fuel

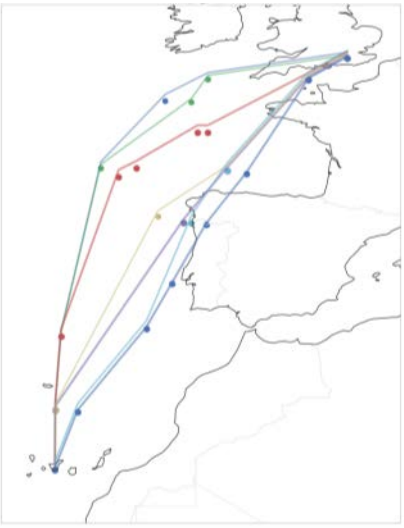

e) Very high fuel

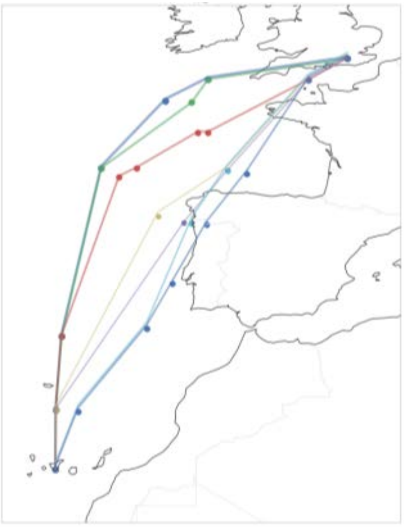

f) FABs level airspace charges

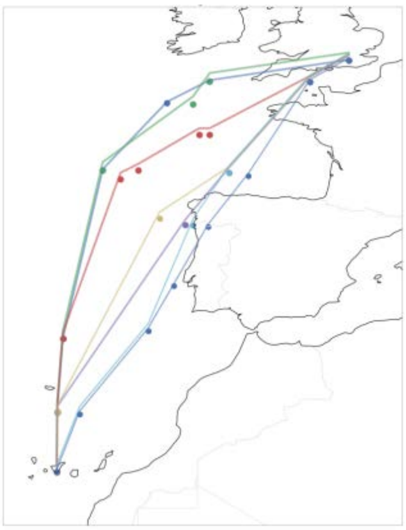

g) Homogeneous airspace charges

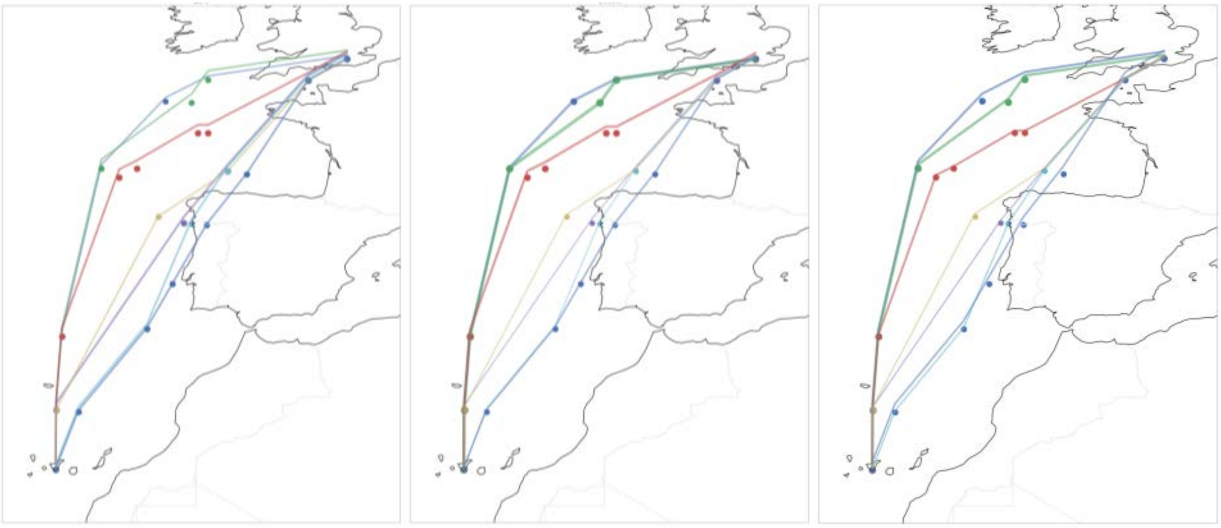

h) Zero cost en-route airspace | i) Double en-route airspace charges | j) Imbalance en-route airspace charges

Figure 124. Routes selected by flights between EGKK to GCTS for different scenarios

Figure 124 shows the different routes between EGKK and GCTS for different scenarios, the width of the line is proportional to the number of flights selecting a route. It is possible to observe that in general changes are note very big, but for example, comparing Figure 124.g (homogenous airspace charges) with Figure 124.i (double en-route airspace charges) it is appreciable that more flights select routes using the Azores and Shanwick airspace in the second case than in the first one.

The variations in demand and in revenues for the different NAS due to fuel prices variations are easier appreciated in Figure 125. See how the demand in the Azores airspace can vary up to more than 20% if the cost of fuel is low or very high, while the Moroccan airspace (GM) can have an increment of 25% if fuel price is very high with respect to the very low case. Also, note that even if some airspaces are unavoidable and have the same number of entry counts (for example Portugal Lisbon (LP)) the amount of revenues obtained might be different as the length of the routes within each airspace differ between scenarios.

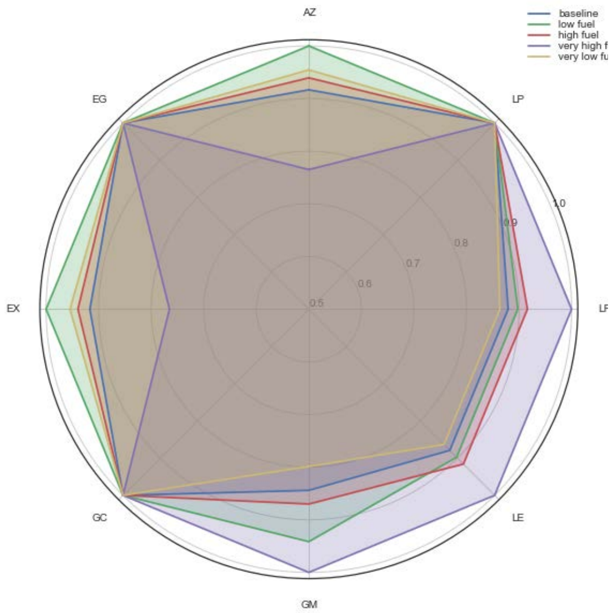

a) changes in demand

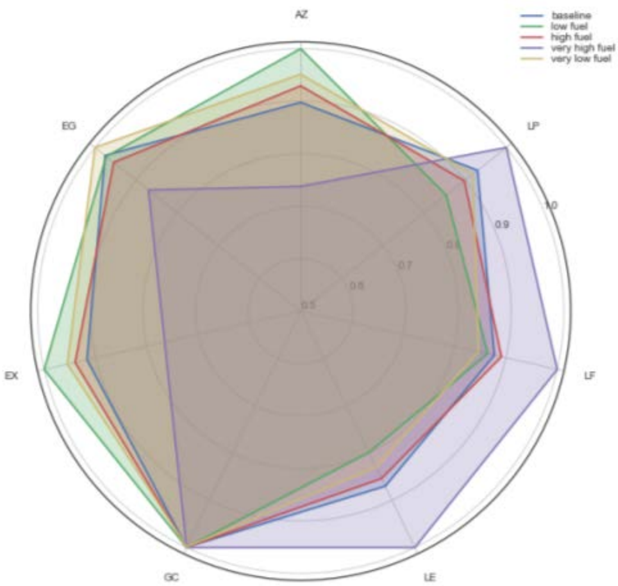

b) changes in revenues

Figure 125. Demand and revenue variations, EGKK–GCTS, when changes in fuel price produced

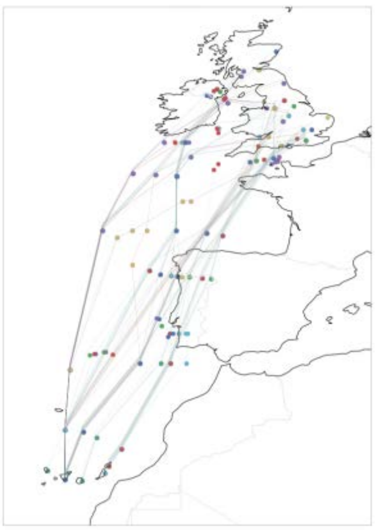

a) Baseline

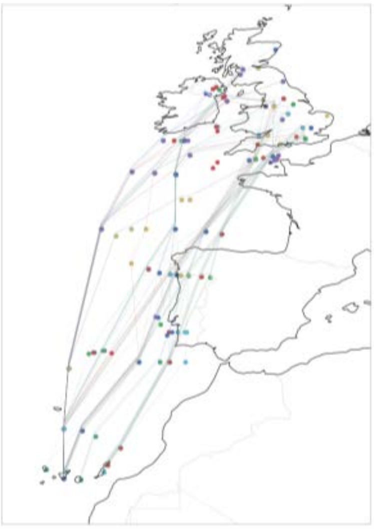

b) Very low fuel

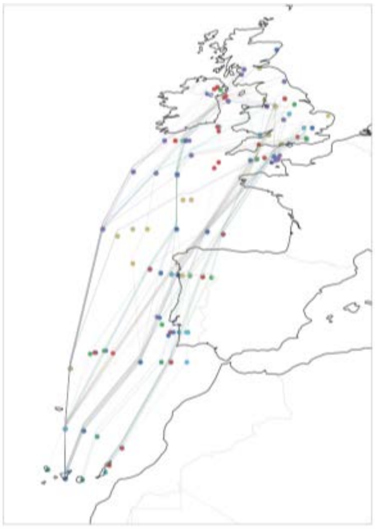

c) Low fuel

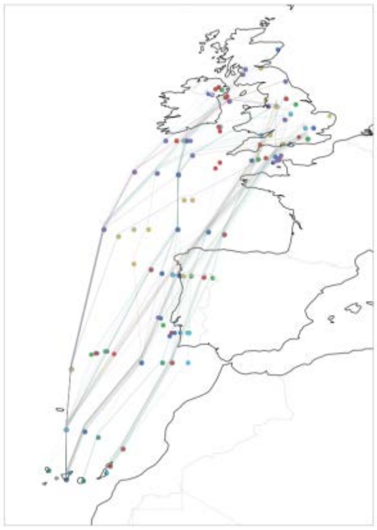

d) High fuel

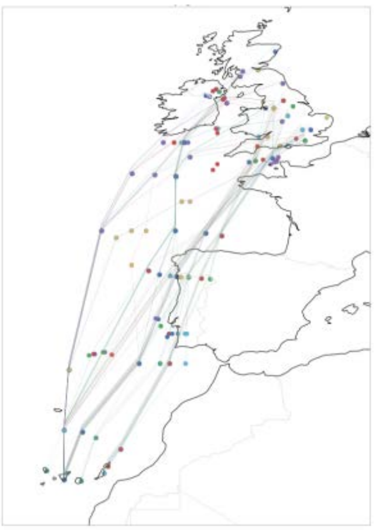

e) Very high fuel

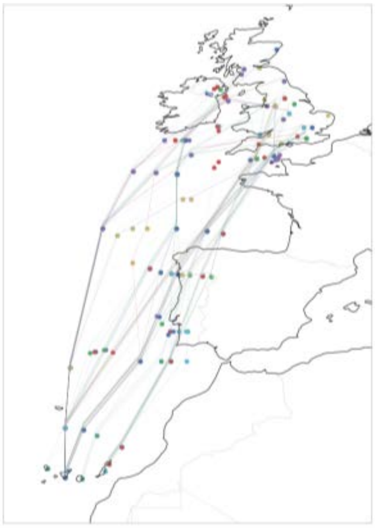

f) FABs

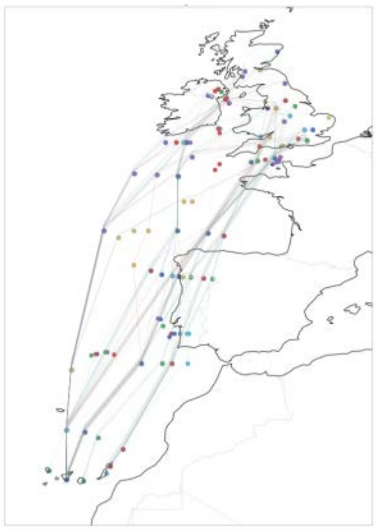

g) Homogeneous

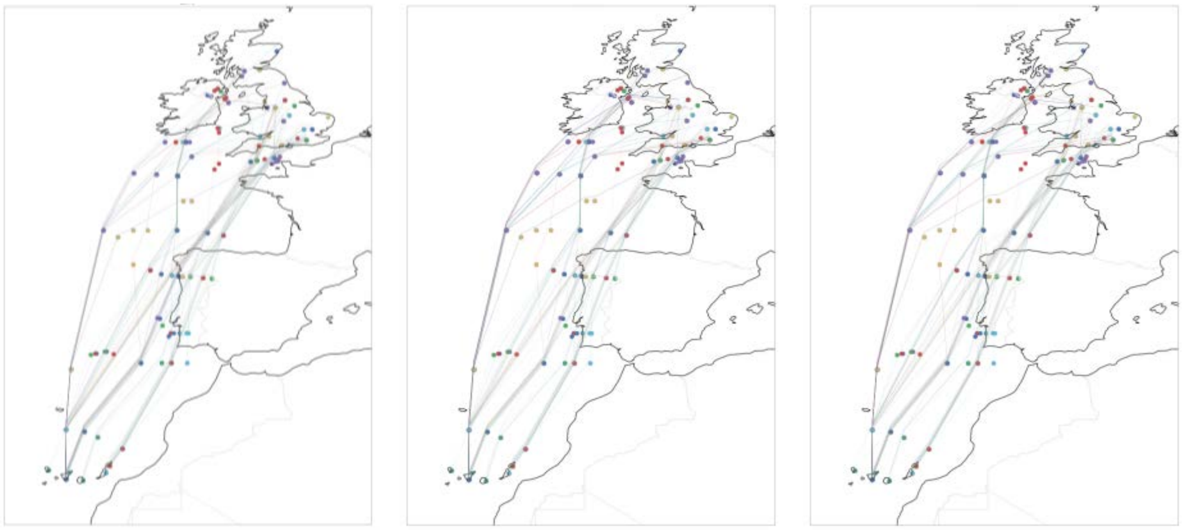

h) Zero i) Double j) Imbalance

Figure 126. Routes selected by flights EG to GC for different scenarios

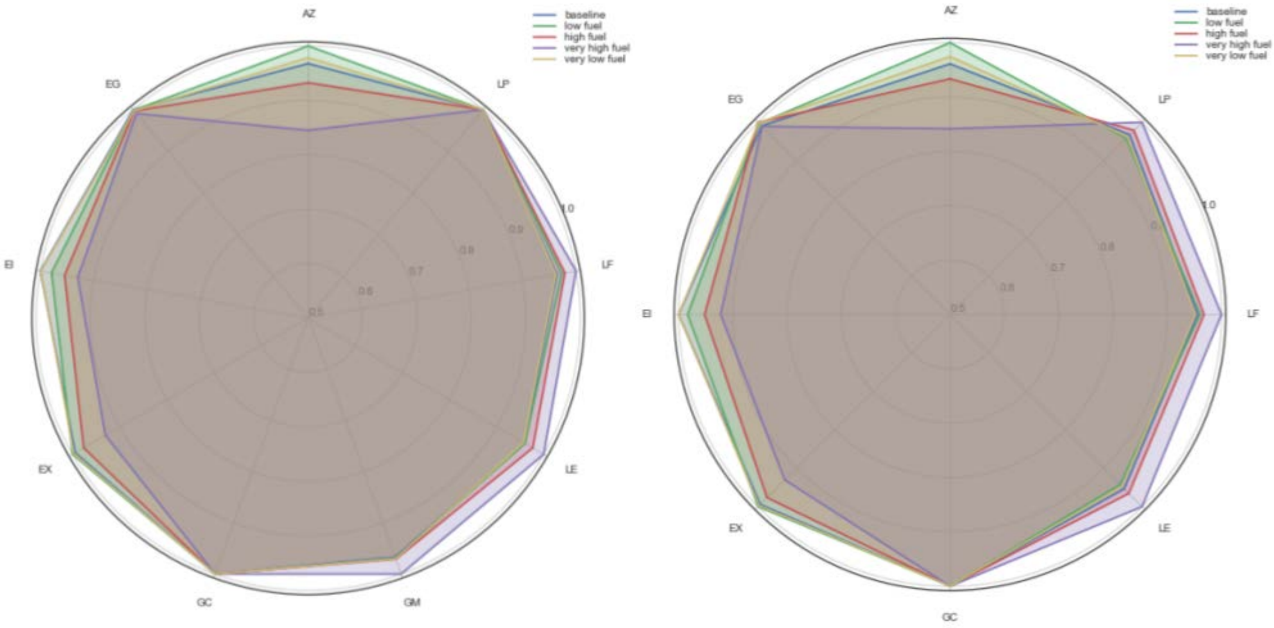

a) changes in demand b) changes in revenues

Figure 127. Demand and revenue variation, EG to GC, when changes in fuel price produced

When all the flights between EG and GC are considered (Figure 126), the variations on demand and revenues due to changes in fuel price are as expected (Figure 127) but in small percentage. This is due to the fact that more flights have less limited options of routes than can avoid more expensive airspaces or that longer routes are required to avoid those airspaces. This shows how this increment of the sample area dilutes the effect of the factor. For example, if all flights between EG and GC are considered the increment in traffic in Morocco (GM) between the very low and the very high fuel scenario is only 5%.

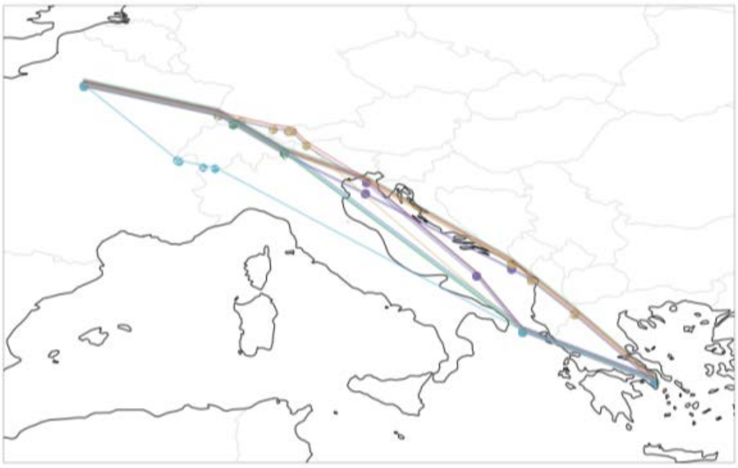

Figure 128. Possible routes between LPFG to LGAV as selected in baseline scenario

Figure 128 presents routes selected for flights between LFPG and LGAV on the baseline scenario. Note that in this case there is the possibility of overflying the Italian airspace or using the cheaper airspace of Croatia, Bosnia and Serbia.

Figure 129 shows the changes in demand for each ANSP for the different fuel scenarios, the influence of fuel is observed on the route preference.

In order to see the impact of changing the airspace en-route charging system and prices, Figure 130 shows side by side, the variations in demand when only the en-route charging scenarios are considered and when these scenarios are executed in conjunction with the fuel variations. Only modifying the airspace charges some variations are appreciated, but when fuel is included the impact is clearer. This indicates that the impact of en-route airspace charges, if existent, is small in comparison with changes on fuel prices.

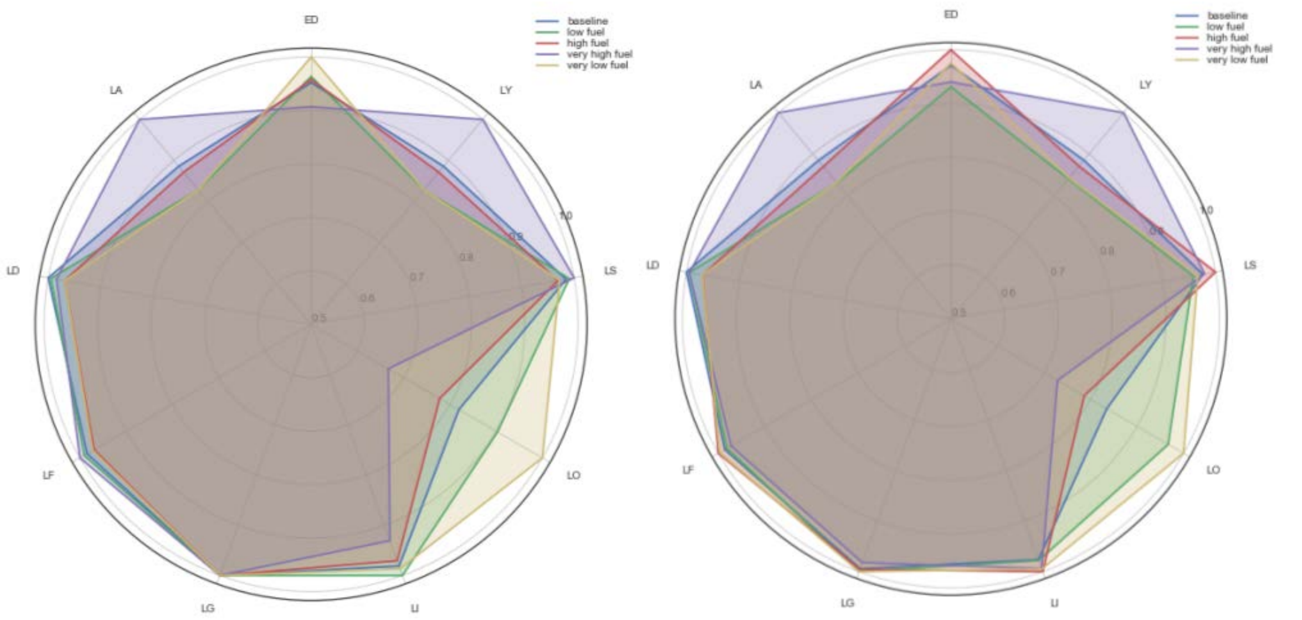

a) changes in demand

b) changes in revenues

Figure 129. Demand and revenue variation, LFPG to LGAV, with changes in fuel price

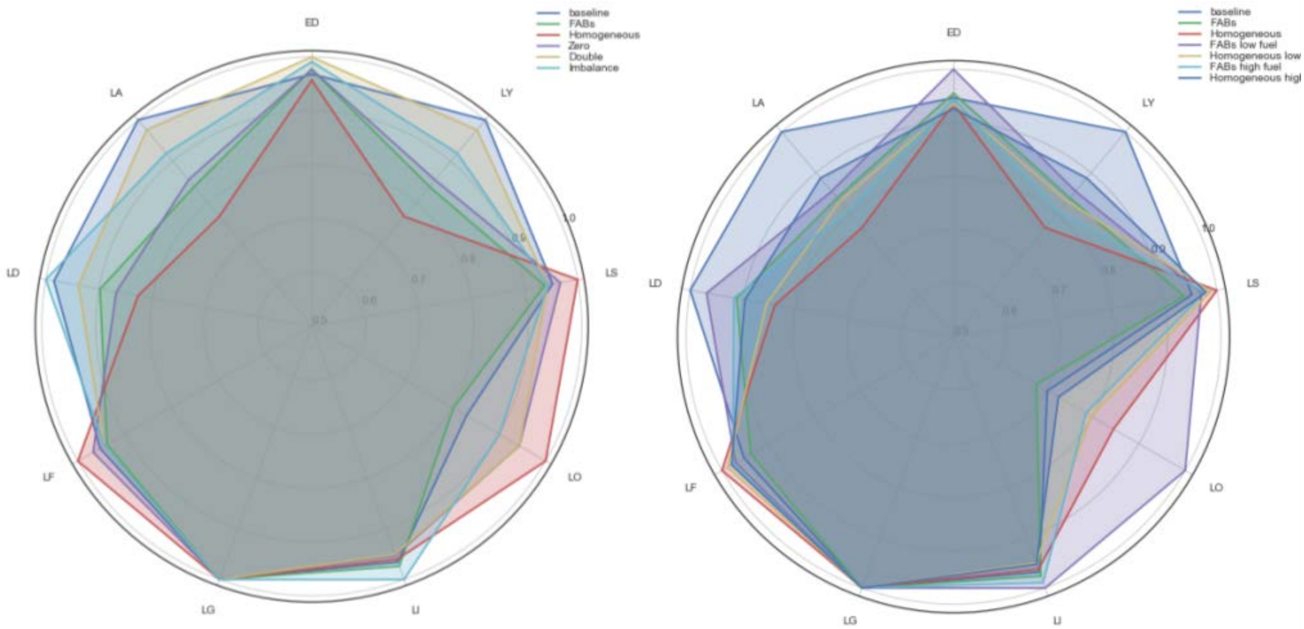

a) charges scenarios

b) charges and fuel scenarios

Figure 130. Demand variation LFPG to LGAV with changes in airspace charges (and fuel prices)

10.2.5 Demand over time by NAS

The final results produced by the pre-tactical layer are temporal variations of demand on the different NAS as aircraft counts in a given time period. Figure 131 shows the entry counts in a 30-minute window for the airspaces of Germany (ED) and Switzerland (LS). As observed there are variations for different scenarios which are linked to the preference of different flight plans. This will be considered by the ATFM Generator to determine the probability of regulating the traffic over the different airspaces.

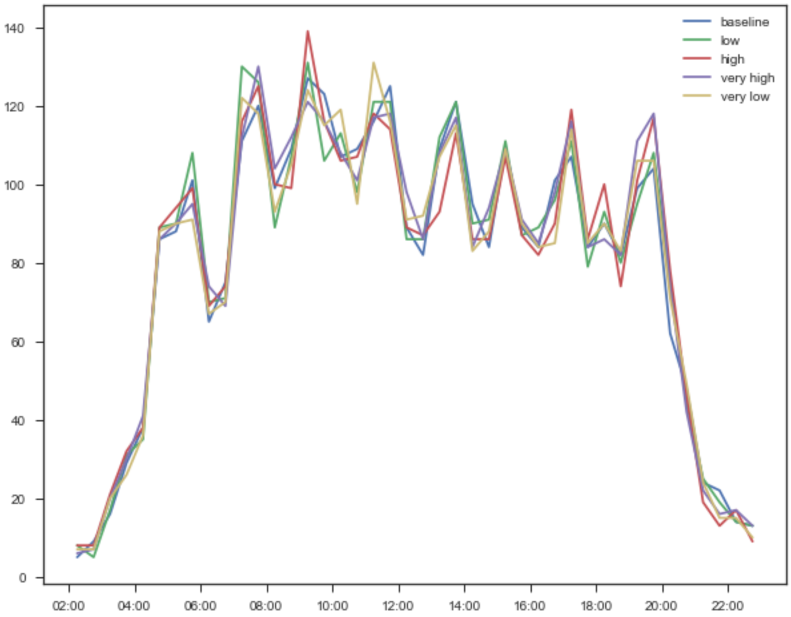

a) ED airspace

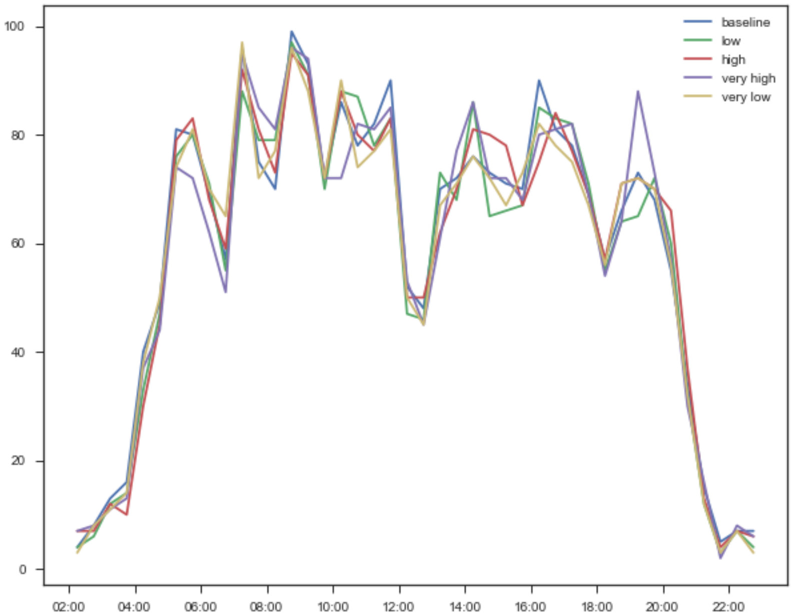

b) LS airspace

Figure 131. Aircraft entries in 30-min time windows by NAS and fuel scenario

-END OF DOCUMENT-