

D5.1 Initial Assessment Report

Deliverable 5.1

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: 07 April 2018
Edition: 01.01.00



Authoring & Approval

Authors of the document

Name/Beneficiary	Position/Title	Date
Luis Delgado / University of Westminster	Consortium Member	08 March 2018
Gérald Gurtner / University of Westminster	Consortium Member	08 March 2018
Samuel Cristóbal / Innaxis	Consortium Member	08 March 2018
Jorge Martín / Innaxis	Consortium Member	08 March 2018
Andrew Cook / University of Westminster	Project Leader	08 March 2018

Reviewers internal to the project

Name/Beneficiary	Position/Title	Date
Graham Tanner / University of Westminster	Consortium Member	09 March 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	10 March 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	10 March 2018	Release	Vista Consortium	New document for review by the SJU
01.01.00	07 April 2018	Release	Vista Consortium	Incorporating 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

This deliverable presents Vista's model and its calibration. The features of each of the model layers (strategic, pre-tactical and tactical) are described along with their calibration. A total of 58 scenarios with four foreground factors are modelled. The results of the layers are produced independently to present the capabilities of the system. These initial results are described and the next steps identified.

Table of Contents

Abstract	3
Executive summary	9
List of Acronyms	10
1 Introduction	11
1.1 Objectives of Vista	11
1.2 Overview of Vista model	11
1.3 Overview of this deliverable	13
1.4 Vista in a few numbers	13
2 Vista Model	14
2.1 Strategic layer	14
2.1.1 Economic Model.....	14
2.1.2 Schedule mapper	22
2.2 Pre-tactical layer	27
2.2.1 Flight plan generator	28
2.2.2 ATFM regulation generator.....	36
2.2.3 Passenger itineraries generator	36
2.3 Tactical layer	36
2.3.1 Purpose	37
2.3.2 Model details.....	37
3 Initial calibration	48
3.1 Strategic layer	48
3.1.1 Economic model.....	48
3.1.2 Schedule mapper	52
3.2 Pre-tactical layer	56
3.2.1 Flight plan generator	56
3.2.2 Air Traffic Flow Management regulation generator	76
3.2.3 Passengers itineraries generator.....	77
3.3 Tactical layer	78
3.3.1 Loading required input data to perform a simulation of all the events.....	79
3.3.2 Calibrating the off-block delay model	80
3.3.3 Passenger reallocation model	80
3.3.4 Simulated taxi times.....	81
3.3.5 Limiting airport capacity through a data driven approach.....	81
3.3.6 En-route time	81
3.3.7 Flights affected by cancellations	81
4 Scenarios and factor modelling	83
4.1 Scenarios modelled in Vista	83

4.2	SESAR-related business factors (BTS)	87
4.3	Strategic layer	88
4.3.1	Economic model.....	88
4.4	Pre-tactical layer	90
4.4.1	Fuel prices (BEO1)	90
4.4.2	CRCO charges (BEO2)	91
4.4.3	Factor and parameters summary	91
4.5	Tactical layer	94
4.5.1	Fuel prices (BEO1)	94
4.5.2	4D trajectory operations (BTS5).....	94
4.5.3	Passenger management tools (BTO4).....	95
4.5.4	Factor and parameters summary	95
5	Initial results	98
5.1	Strategic layer	98
5.1.1	Evolution of baselines in 2035 and 2050.....	99
5.1.2	Impact of fuel price	104
5.1.3	Impact of 4D trajectories	106
5.1.4	High-level view for trade-off analysis.....	109
5.2	Pre-tactical layer	112
5.2.1	Impact of fuel cost on flight plan selection	113
5.2.2	Impact of en-route airspace charges on flight plan selection	115
5.2.3	Impact of fuel cost on revenue per ANSP	115
5.2.4	Individual origin destination examples	116
5.2.5	Demand over time by NAS	127
5.3	Tactical layer	128
5.3.1	Isolated factor results.....	129
5.3.2	Cross-factor results	129
6	Next steps and look ahead	131
6.1	Model layer integration and consultation in progress	131
6.2	Features to be potentially deployed in final model version	131
7	References	132
8	Annex: Interpolated and extrapolated targets	133

Table of Tables

Table 1. Tactical model, flight state variables	37
Table 2. Tactical model, flight derived simulation times	38
Table 3. Tactical model, passenger itineraries variables.....	39
Table 4. Tactical model, passenger derived simulation times	40
Table 5. Tactical model, passenger wait variables	44
Table 6. Tactical model, passenger connecting time attributes	44
Table 7. (Tactical) Parent airports table.....	45
Table 8. (Tactical) Nested airports table	45
Table 9. Tactical model, passenger connecting time attributes	46
Table 10. Tactical model, airports waiting passengers table	46
Table 11. Tactical model, taxi time table.....	46
Table 12. Tactical model, airport runway table.....	47
Table 13. Quantiles and mean for FP distance (NM).....	67
Table 14. Quantiles and mean for FP taxi times and buffers (mins).....	69
Table 15. Passenger itinerary data and group model.....	79
Table 16. Flight schedule mapping.....	79
Table 17. Scenarios simulated by Vista.....	84
Table 18. Correspondence between background scenarios and operation types	88
Table 19. Impact of SESAR KPIs on inner parameters of the economic model	89
Table 20. GDP growth in the different scenarios	90
Table 21. Effect of other factors on inner parameters of the model.....	90
Table 22. Summary of factors modelled in pre-tactical layer	92
Table 23. Summary of factors modelled in tactical layer	96
Table 24. Tactical metrics computed and baseline reference value.....	128
Table 25. Time-based operations, SESAR targets used for factors	133
Table 26. Trajectory-based operations, SESAR targets used for factors	134
Table 27. Performance-based operations, SESAR targets used for factors	135

Table of Figures

Figure 1. Overview of Vista model	12
Figure 2. High-level view of the strategic layer	14
Figure 3. Elementary taxons	24
Figure 4. Composite taxons	24
Figure 5. High level pre-tactical layer blocks	27
Figure 6. High level pre-tactical layer blocks	28
Figure 7. Example of route generation LIMF – LFBT	30
Figure 8. Example of route generation MUHA – LSZH	31
Figure 9. From flight schedules to routes per flight schedule	32
Figure 10. From routes to trajectories	33
Figure 11. Regions for which the air navigation charges are computed	35
Figure 12. Flight times and buffers	36
Figure 13. Tactical layer event chain diagram	40
Figure 14. Airline cost data comparison between empirical data and outputs of the model	52
Figure 15. Number of tail numbers (left) and flights (right) by closed or open loops	53
Figure 16. Distribution of elementary taxons, characterised by number of airports (abscissa)	54
Figure 17. Distributions of the most important (17 first) combined taxons	54
Figure 18. Departure times of aircraft (from first airport) for taxons (1, 0, 0), (2, 0, 0), and (7, 0, 0)	55
Figure 19. Distribution of turnaround times (in hours) for taxons (1, 0, 0), (2, 0, 0), and (7, 0, 0)	56
Figure 20. Example of route clustering EBBR – LTBA	58
Figure 21. Example of route clustering DGAA – EGLL	59
Figure 22. Example of flight level, speed request	59
Figure 23. Example average cruise flight level	60
Figure 24. Example distributions of average cruise FL requested for B737 by FP distance	61
Figure 25. Example of distributions of speed requested for A320 by flight level	62
Figure 26. Example of estimated climb and descent phase	62
Figure 27. Example of distributions of climb and descent distance for B737 by FP distance	63
Figure 28. Example of linear fitting of climb and times for A319	64
Figure 29. Examples of average wind distributions between NAS	65
Figure 30. Flight plan distances: historical flights and average for FP selected per schedule	66
Figure 31. Histogram of flight plan distance for all FPs generated	66
Figure 32. Cumulative probabilities of FP distance (NM)	67
Figure 33. Taxi times and schedule buffers for historical and selected FPs	68
Figure 34. Cumulative probability of taxi and buffer time	69
Figure 35. Taxi and buffer times as a function of FP distance (NM)	70
Figure 36. Fuel as a function of FP distance (NM) per aircraft type	72
Figure 37. En-route air service cost as a function of FP distance (NM) per aircraft type	74
Figure 38. En-route air services revenue per NAS	74
Figure 39. En-route revenues per NAS from FPs generated and reported service units	75
Figure 40. En-route service units per NAS from FPs generated and reported service units	76
Figure 41. Probability distribution of ATFM delay	77
Figure 42. Average daily cancellations 2014	82

Figure 43. Monthly share of operational cancellations 2014.....	82
Figure 44. Six key metrics for different stakeholders in the baseline scenarios	100
Figure 45. Airport metrics in the baseline scenarios	101
Figure 46. Airline metrics in baseline scenarios	102
Figure 47. ANSP metrics in baseline scenarios	102
Figure 48. Pax metrics in baselines scenarios	103
Figure 49. Airline metrics by scenario	104
Figure 50. Operational cost structure of the airlines by scenario (percentages)	105
Figure 51. Further airline metrics by scenario	106
Figure 52: Airline metrics under different scenarios, testing 4DT implementation levels	107
Figure 53. Testing 4DT implementation levels for different scenarios with the H35 baseline.....	108
Figure 54. Time evolution of key metrics when testing 4DT implementation levels	109
Figure 55. 'Radar' plots for the 5 KPIs in the baseline scenarios	110
Figure 56. 'Radar' plots for the 5 KPIs when changing price of fuel in 2050	111
Figure 57. 'Radar' plots for the 5 KPIs in when changing the 4DT implementation level in 2050...	112
Figure 58. Impact of fuel cost on FP selection characteristics	113
Figure 59. Impact of fuel cost on FP selection characteristics; flights with one route option.....	114
Figure 60. Impact of fuel cost on FP selection characteristics; flights with >1 route option	114
Figure 61. Impact of en-route charges on flight plan selection characteristics.....	115
Figure 62. En-route airspace usage revenues per NAS for different fuel scenarios.....	116
Figure 63. Impact of fuel cost on FP selection characteristics for GCTS to EGKK flights	117
Figure 64. Impact of fuel cost on FP selection characteristics for all EG–GC flights.....	117
Figure 65. Impact of en-route charges on FP selection for GCTS to EGKK flights.....	118
Figure 66. Impact of en-route charges on FP selection for all EG–GC flights	119
Figure 67. Routes selected by flights between EGKK to GCTS for different scenarios.....	121
Figure 68. Demand and revenue variations, EGKK–GCTS, when changes in fuel price produced ...	122
Figure 69. Routes selected by flights EG to GC for different scenarios	124
Figure 70. Demand and revenue variation, EG to GC, when changes in fuel price produced	124
Figure 71. Possible routes between LPFG to LGAV as selected in baseline scenario	125
Figure 72. Demand and revenue variation, LPFG to LGAV, with changes in fuel price.....	126
Figure 73. Demand variation LPFG to LGAV with changes in airspace charges (and fuel prices)	126
Figure 74. Aircraft entries in 30-min time windows by NAS and fuel scenario	127
Figure 75. Initial results of the tactical layer for isolated factors	129
Figure 76. Initial results of the tactical layer for paired factors.....	130
Figure 77. Initial results of the tactical layer when the three factors simultaneously active	130

Executive summary

This deliverable presents the first main batch of results obtained from the model, together with a detailed description of all the parts of the model, the analyses performed to prepare the data, and the initial calibration of the different parts of the model.

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 the 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.

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 model, economic model and a schedule mapper to generate demand and capacity for the different stakeholders, flight schedules and passenger flows. The pre-tactical layer transforms the output of the strategic layer into individual flight plans, passenger itineraries and ATFM regulations delay. Finally, the tactical layer runs the day of operations, tracking flights and passengers and reacting to the tactical situation in the system.

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 so far are described in Section 4, with the initial results described layer by layer in Section 5. At the current state of development of the model, Vista is able to show how different factors affect KPIs on the different layers independently.

A consultation will be carried out to obtain feedback on the model assumptions and on the scenarios so that the different model blocks can be developed and finally integrated into a single model that will produce the final results, which will be presented in D5.2.

List of Acronyms

4DT	4D trajectory	ETA	Estimated Time of Arrival
ACI	Airport Council International	ETOT	Estimated Take Off Time
AIBT	Actual In Block Time	FAB	Functional airspace block
AIRAC	Aeronautical Information Regulation And Control	FL	Flight level
ANSP	Air Navigation Service Provider	FP(G)	Flight Plan (Generator)
AO	Aircraft operator	FSC	Full service carrier
AOBT	Actual Off Block Time	GCD	Great circle distance
ATA	Actual Time of Arrival	GDP	Gross domestic product
ATC	Air Traffic Control	GDS	Global Distribution System
ATFM	Air Traffic Flow Management	IATA	International Air Traffic Association
ATM	Air Traffic Management	KDE	Kernel Density Estimation
BADA	Base of Aircraft Data	KPI	Key performance indicator
BEO1	Business Other Economic Factor 1: Price of fuel	LCC	Low cost carrier
BEO2	Business Other Economic Factor 2: Airspace charges	MER	Mercury (tactical layer of Vista)
BTO4	Business Technology Others Factor 4: Passenger reaccommodation tools	MTOW	Maximum take-off weight
BTS5	Business Technology SESAR Factor 5: 4D Trajectory Management	NAS	National Airspace System
CHT	Charter carrier	NM	Nautical miles
CODA	Central Office for Delay Analysis	NOP	Network Operations Plan
CRCO	Central Route Charges Office	PRU	Performance Review Unit
DDR	Data Demand Repository	REG	Regional carrier
ECAC	European Civil Aviation Conference	SIBT	Scheduled In Block Time
EIBT	Estimated In Block Time	SOBT	Scheduled Off Block Time
EM	Economic Model	UTC	Universal Time Coordinated
EOBT	Estimated Off Block Time	WID	World Wealth & Income Database

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 agents, 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 passenger level, tracking the evolution of delay, passenger connections and the tactical decisions carried out by the stakeholders. 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 mostly 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.

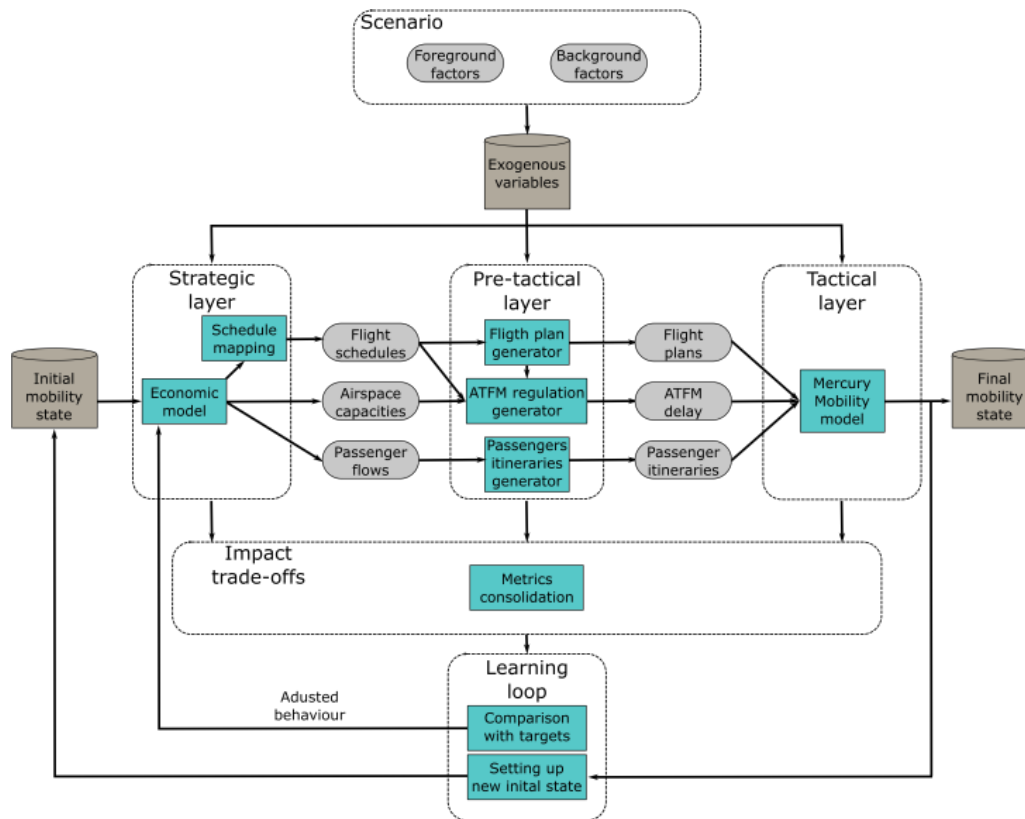


Figure 1. Overview of Vista model

So far, the project has focused on the core layers of the model: strategic, pre-tactical, and tactical. The learning loop displayed in Figure 1 has not been implemented yet. As highlighted in Deliverable 3.1, its implementation was uncertain due to the already intensive effort needed for the core capabilities. It is now foreseen that the team will not develop this learning loop, which would not profoundly change the output of the model in any case. Regarding the impact trade-off assessment layer, its development has just started, with the first results now being produced. It will be relatively less demanding in any case, since it will include a few tools to perform useful data analyses. Its output will be presented in the final deliverable.

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. This information could be used by a learning loop to improve the decision-making process of the agents at the strategic level.

1.3 Overview of this deliverable

This deliverable presents the initial assessment of some scenarios. Results of the different layers are independently presented to show the capabilities of the layers to capture the impact of some factors. The deliverable presents:

- the details of the modelling for each of the layers;
- a significant amount of effort has been used to ensure that results obtained by the model are realistic; this is achieved through the calibration of the model; the calibration of each layer is presented in this deliverable;
- to test the capabilities and validate the results obtained by each layer, different factors are modelled creating a set of scenarios. This is described in Section 4.
- the initial results obtained are described for the different scenarios providing insights into the impact of the factors on different metrics;
- the next steps required to complete the model development and integration, and to test new scenarios, are described.

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 6.5 GB of input/output data;
- A GitHub repository with 303 commits;
- 30 975 lines of code;
- 56 scenarios simulated so far;
- Around 20 factors simulated so far.

The opinions expressed herein reflect the authors' views only. Under no circumstances shall the SESAR Joint Undertaking be responsible for any use that may be made of the information contained herein.

2 Vista Model

2.1 Strategic layer

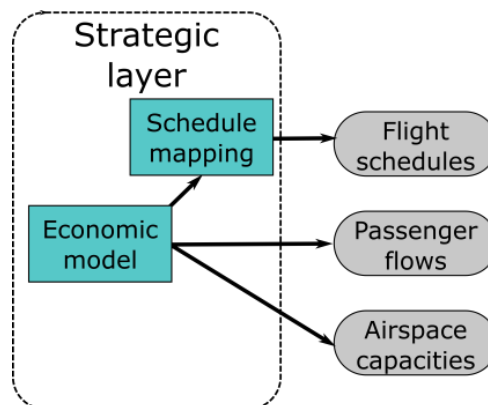


Figure 2. High-level view of the strategic layer

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.

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.
- 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 connection**). Notionally represents all the passengers going from O to D through any legitimate itineraries (sequence of airlines and sequence 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. 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. Here is a short list of what happens in the models:

- Based on the scenario, parameters of 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 their unit rate so that their profit is null.
- Airports predict the traffic for this round. They compute if they should increase their capacity, based on their cost function.
- Airline compute if they should open some new routes, based on expected demand and their cost function.
- 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.
- New flights are created, construction for extra capacity at airports is initiated.
- 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:

- δt_O and P_O , respectively the delay and aeronautical charges of the departure airport, from the corresponding agent,
- δt_D and P_D , respectively the delay and aeronautical charges of the arrival airport, from the corresponding agent,

- δt_{ATFM}^j and u_j , respectively the ATFM delay generated by each ANSP crossed by the flight plan, and its unit rates, from the ANSP agents,
- the fuel price p_f , from the 'world'.

It then computes the following cost for the flight plan i , with total distance d_i :

$$c_0^i = \chi \left(\delta t_O + \delta t_D + \sum_{j \in \text{ANSPs}} \delta t_{ATFM}^j \right) + (a + b d_i + c d_i^2) p_f + \sum_{j \in \text{ANSPs}} \left(u_j \frac{d_i^j}{100} \sqrt{\frac{\text{MTOW}}{50}} \right) + P_O + P_D + (c_{\text{crew}} + c_{\text{maint}}) d_i$$

where χ is the cost of delay of the airline, a , b and c 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, d_i^j is the distance flown in ANSP j airspace and c_{crew} and c_{maint} are the cost for crew and maintenance, computed by kilometre flown. Once the cost of each potential flight has been computed, the flight weights the flight plans according to a logit rule:

$$p_i = \frac{e^{-c_0^i/s}}{\sum_j e^{-c_0^j/s}}$$

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:

$$c_o = \sum_i p_i c_0^i$$

Note that a negative term is sometimes added to the cost to ensure that the flight is profitable at least in the initial state*.

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:

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

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:

$$c(S) = c_o S + c_c S^2$$

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.

(ii) Airlines

The airline agents are mostly used to create new flights. In order to open a new connection, i.e. a new flight agent, the airline examine whether it would be profitable for them. For now, the airline only considers a new route if:

- at least one other airline operates the same route,
- the airline can make a connection with some other flights, i.e. the new route must be connected to the existing network of the airline.

It first computes all the itineraries which would be open to them with this new route, including direct and hub-connecting routes*. Then they proceed to estimate the price and demand on these routes. In order to do this, they compute the average price already paid by passengers on these routes with other airlines, disaggregating them to keep only the relevant leg, by subtracting the price of the other legs. The expected demand is simply given by the same passengers going through this leg.

It then moves to estimate the cost of operating such a route. In order to do this, it uses a similar cost function than the existing flights themselves, averaging the relevant coefficients when needed (cost of crew, maintenance, and fuel consumptions. If the expected optimal (profitable) supply is greater than a given threshold*, the new flight is created.

(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:

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

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 $w = \lambda / (2\mu(\mu - \lambda))$. 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 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 pax 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 ref. [1] 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. After some search, making such a list proved very difficult, so this feature is dormant for now in the model. This is one of the features that are on the list of possible improvements (see Section 6.1).

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

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’, except that we here chose an exponential law:

$$\delta t = A e^{\frac{T}{c}}$$

The choice of an exponential is driven by partially by the fact that exponentials fit slightly better the data than a linear law, even if the fits are relatively poor. For more details on how we set the

parameter A , please refer to the calibration*. Given the target delay and having predicted a level of traffic (in the same way that the flights predict price), inverting this equation thus gives us the capacity needed for the next round to be under the target delay. Once this is fixed, the ANSP computes the corresponding cost c and fixes the unit rate u , to have zero profit:

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

Where V is the term of revenue corresponding to $\sum_i d_i \sqrt{MTOW_i/50}$ where the sum is 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 ANSPs' behaviours have been set to match current behaviour. In the future, it is foreseen that the ANSPs will undergo some large changes, albeit with details as yet unknown. It is planned for the final deliverable to test at least one different set of rules for the ANSPs, and it is listed as one of the possible features for next version (see Section 6.1).

(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 α 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 than 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 that 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**, they 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 on. 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 a 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 available taxons (see next subsection);
- Compute likely turnaround times per taxon;
- Compute the likely departure 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.

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. It then tries to find the best patterns, and finally the departure times.

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 $(1, 0, 0) \rightarrow (2, 0, 0)$. 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.

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 $DD = D - S$, 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 $DD_i > 0$, i.e. where there is more demand than the supply, and note them DD^* and P^* . The cost is then:

$$c = c_c N_p + DD^* \cdot P^*$$

The first term is a cost of capital, where c_c is the cost for creating a pattern, loosely corresponding to the cost of an aircraft (buying or renting). 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

Once useless patterns have been removed, each airline uses a two-stage stochastic optimisation to add patterns. It begins by considering a 'chromosome', a vector encoding how many of each taxon is added to the airline network. The airline tries to randomly modify this chromosome to find the best set of taxons, using a simulated annealing procedure (first stage). The random tries are driven by the empirical distributions of taxons, extracted from data. This means that in practice that the allele corresponding to the taxon (1, 0, 0) is vastly more probably chosen and modified during the procedure than any other, as it is the main building block of the airlines networks.

For each configuration of the chromosome (i.e. numbers of taxons (1, 0, 0), numbers of taxons (2, 0, 0, etc), the airline has to estimate a cost for it. In order to do this, it needs to convert the chromosome into actual patterns and derive a cost based on the master cost function presented above. To do this, the airline uses a second simulated annealing procedure (second stage). Given a set of taxons, it tries to randomly assign airports to them using historical data on times of departure and travelling times.

Specifically, for a given taxon, the algorithm starts by choosing the leg with the highest difference between demand and supply. It then starts to test several airports for the intermediate ones, using travelling times and turnaround times as computed from data. In particular, the turnaround times are sampled from the actual turnaround times for a given taxon. This ensures than other hidden variables are taken into account in the procedure. For instance, a taxon (1, 0, 0) should probably correspond to a long flight (given that the aircraft makes a return in 72 hours), and thus takes implicitly into account crew needs as well as aircraft maintenance by having a higher chance of a long turnaround. The algorithm also tests several aircraft type to be applied on the pattern allowing different numbers of seats n_p .

The cost for a given pattern is then computed by using a function derived from the master function:

$$c = p_d + c_c - \sum_{i \in \{j | DD_j > 0\}} \min(n_p, DD_i) P_i$$

The third term is almost the same as the second one in the master function, except that we capped the gain by capacity of the aircraft selected, n_p . The sum is carried over all the legs in the pattern where the demand is superior to the supply. The first term is a penalising one, equal to 0 if the total duration of the flight is smaller than the time-window (72h), and typically very high otherwise. It ensures that the solution is compliant with our assumptions regarding the time window. Note that the minus sign in front of the third term comes from the fact that here we are testing if the added capacity n_p on leg i is decreasing the shortfall, whereas in the master equation above we were computing the shortfall after patterns were applied.

Once the cost of a pattern is computed the algorithm comes back to trying several of them choosing the best one for each taxon. It can then test different number of taxons (changing the chromosome described above), which finally leads to an 'optimal'. Note that the optimality of the solution cannot be tested. As with any other heuristic, the solution of a simulated annealing procedure is never guaranteed to be the best solutions (here the minimum cost for the airline). However, it is known to perform very well in a wide range of problems. Moreover, from a realistic point of view, we are not interested in the perfectly optimal solution. First, because we do not have access to all the parameters in play for the airline, secondly we do not know for sure what is the objective function of the airlines, and finally because in any case it is likely that the airlines make suboptimal choices anyway.

2.2 Pre-tactical layer

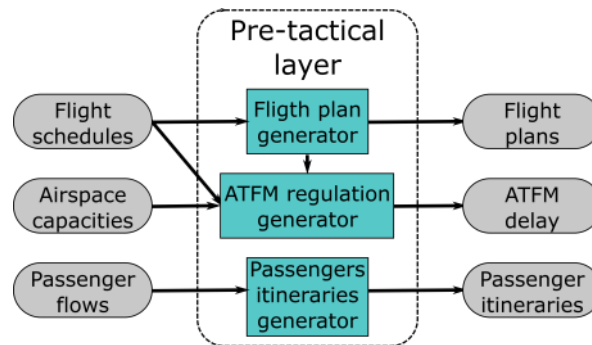


Figure 5. High level pre-tactical layer blocks

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;
- ATM 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)).

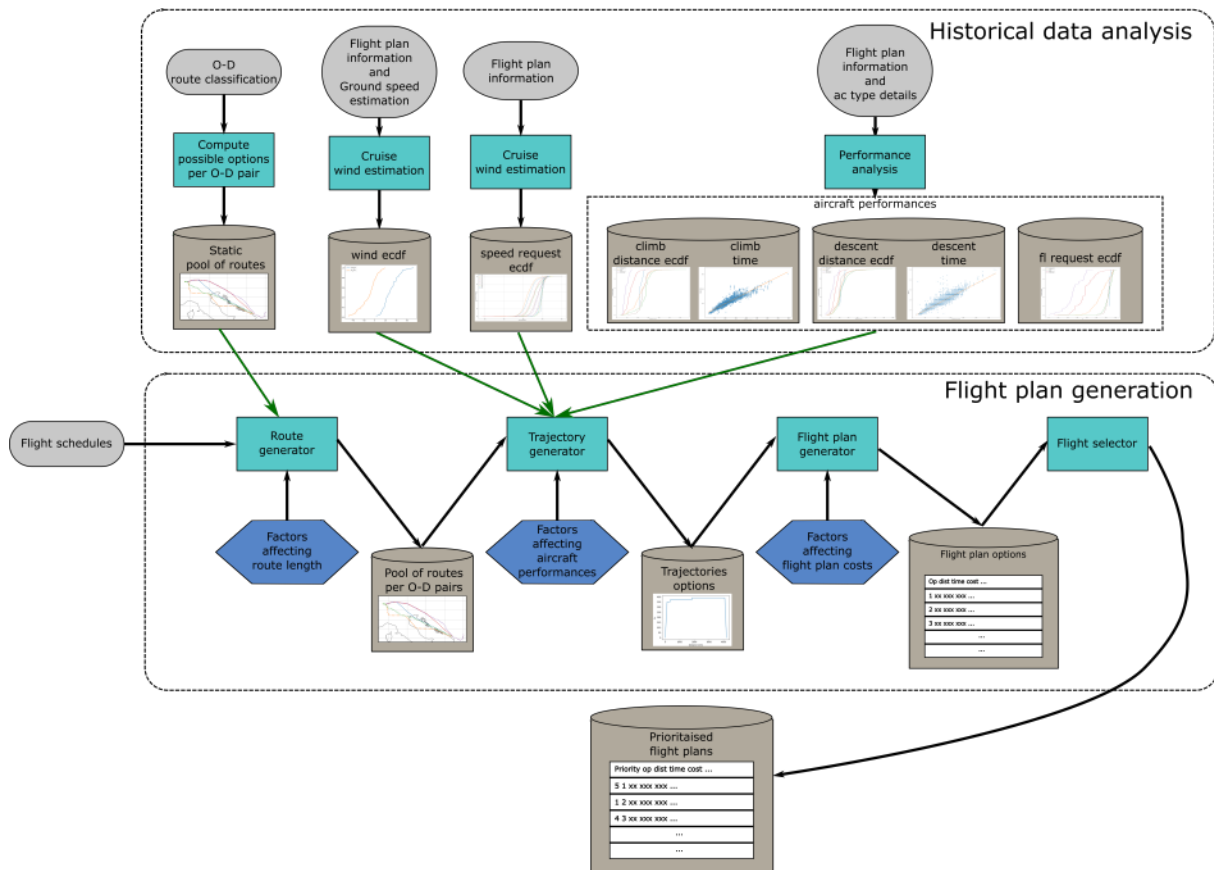


Figure 6. High level pre-tactical layer blocks

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.

2.2.1.1 Route generator

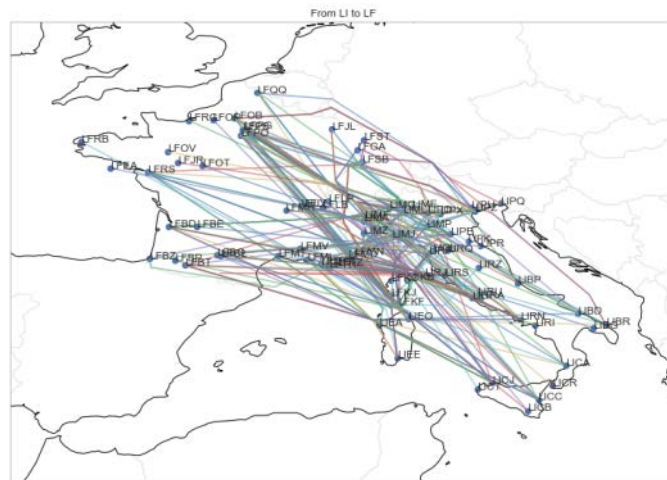
For each scheduled origin-destination pair a set of possible routes are computed. These possible routes are based on the historical possible routes flown between these origins and destinations (see Section 3.2.1.2 for details on how these options are computed from historical data). If the origin-destination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

All the routes between the NAS¹ 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.

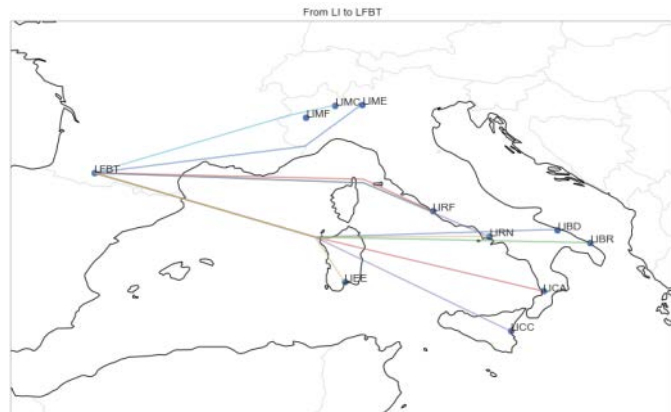
¹ Note that if there are no flights between a given origin and destination, all flights between their NAS of origin and destination are considered.



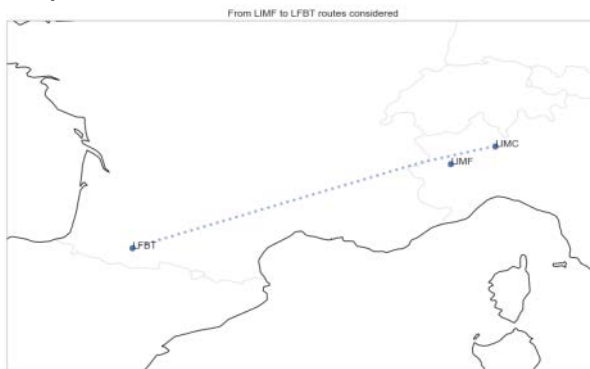
a) All possible routes between LI and LF



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.

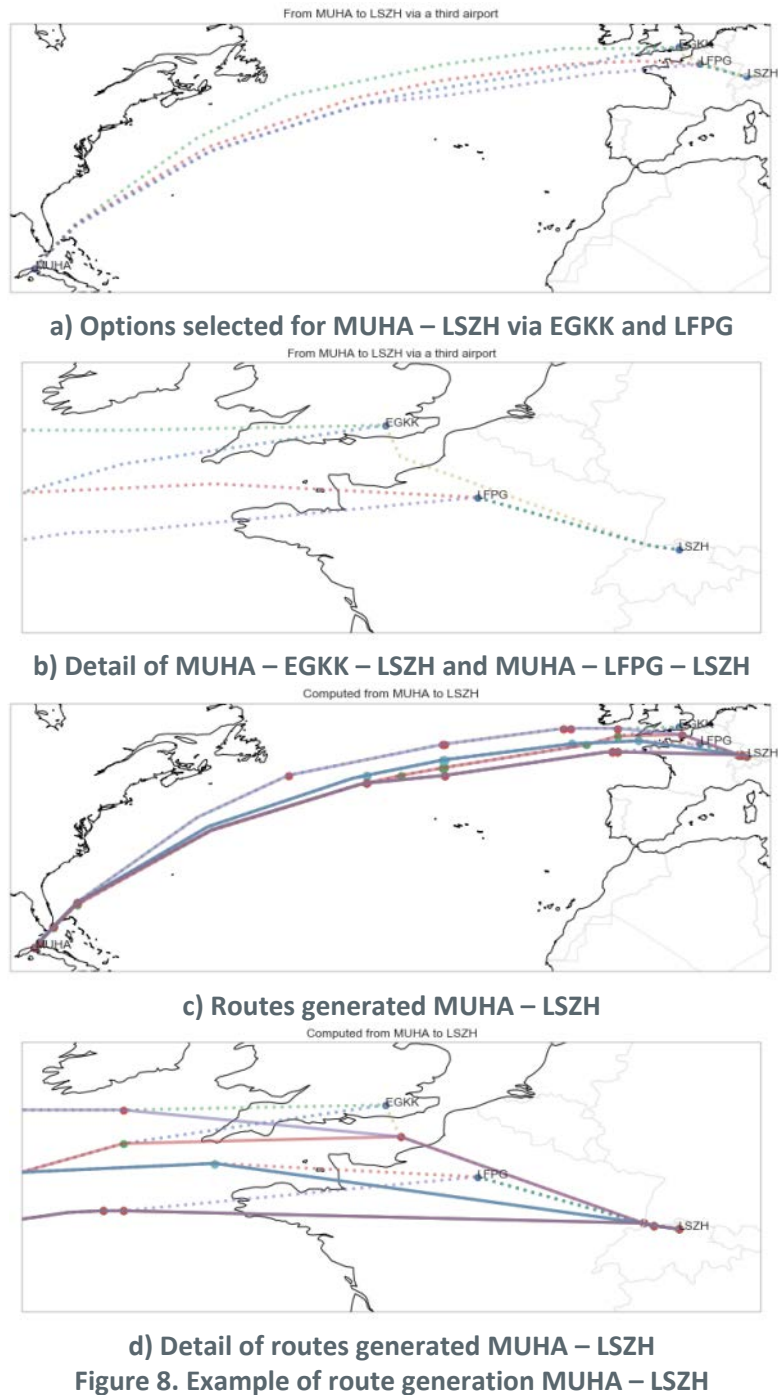


Figure 8. Example of route generation MUHA – LSZH

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:

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

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, for example by reducing the distance within each NAS if free routes are implemented.

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.

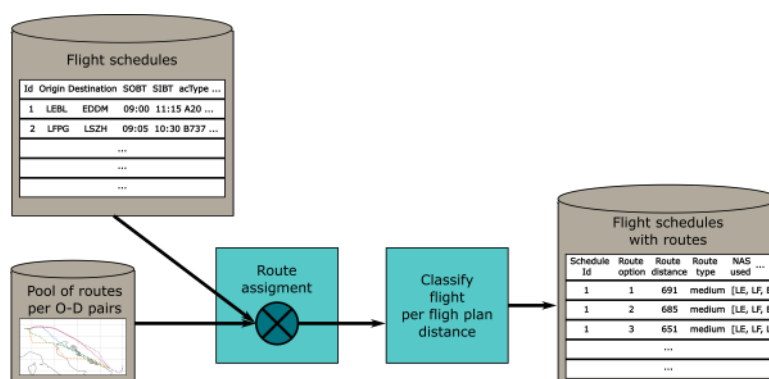


Figure 9. From flight schedules to routes per flight schedule

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.

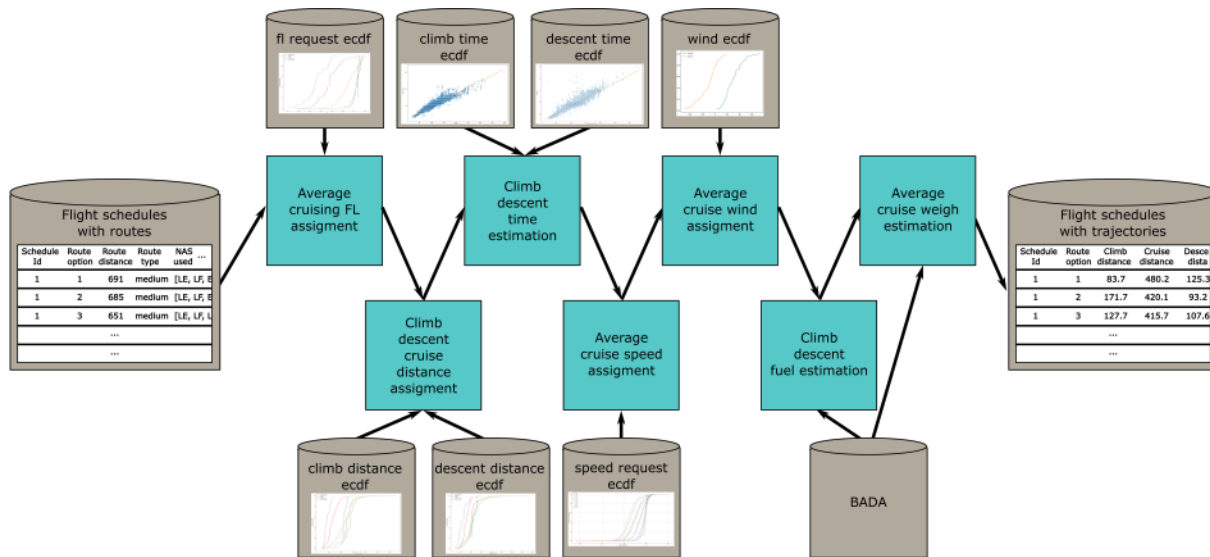


Figure 10. From routes to trajectories

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 performance, 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.

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.

If there are factors that affect the aircraft performances and hence the trajectories, they will be considered in this block. In the current version of the model, the trajectories generation are based on the historical analysis of trajectories and on the aircraft performance models of BADA.

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 rate 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.

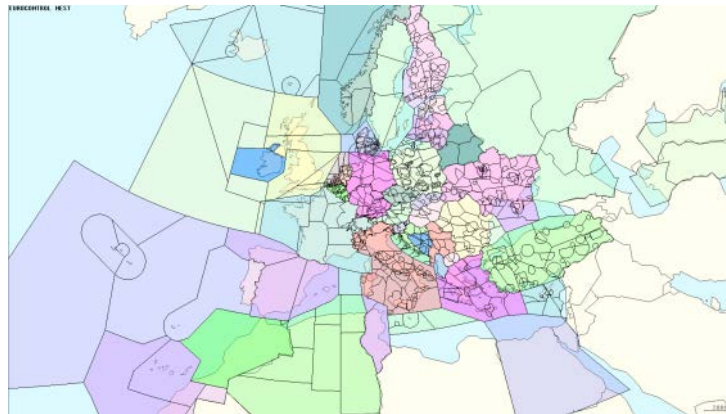


Figure 11. Regions for which the air navigation charges are computed

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).

There are factors that affect the cost of fuel and how air traffic charges are computed. These will be modelled in this block and affect the operating cost of each possible trajectory.

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.

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 and airspace cost but also 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

In the current version of the model, a logit rule, which considers only the cost of fuel and the cost of airspace charges, 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.

In the current implementation of Vista, this is done with an experimental cumulative probability distribution function computed from the analysis of the ATFM regulations of a year (AIRACs 1313-1413). There is a probability of having a regulation which is the most penalising one for the flight and in that case the delay assigned follows the cumulative probability distribution function.

The model is able to compute the demand of the number of flights inside each NAS over time considering that the flights are selecting the flight plan with the highest priority from the Flight Plan Selector. This demand will be used to modify the probability of having delay assigned as more demanded regions will experience higher ATFM delay. On the other hand, a statistical analysis of ATFM regulations due to other factors (such as related to weather) will be performed and hence the delay assigned will be a combination of the probability of being affected by other factors regulations (such as weather) and of being affected due to congestion (high demand in regions/airports). This variation of ATFM delay as a function of demand and cause will be incorporated in D5.2.

2.2.3 Passenger itineraries generator

The current version of Vista does not implement this block yet. The itineraries used in the tactical layer are those computed for the current (SEP14) scenario. In the next (final) version of the model, the passenger flows will be transformed into individual passenger itineraries considering the schedule and the minimum connecting time for the passengers at the airports.

2.3 Tactical layer

One of the layers of the Vista model is the Tactical layer. This layer includes stochastic components and feedback loops. The model takes into account different stakeholders for capacity and delay management. The model is written in Python from scratch using statistical, analytical and simulation packages. A single simulation run could take between 5 and 20 minutes to generate tactical metrics using a mid-performance workstation.

The model is based on a legacy model from previous work, written in MATLAB, that was hard to extend and maintain. The ad-hoc event based scheduled presented in that model is replaced in the current model with a discrete event based scheduled included in the library. Some other models are simplified from the legacy model to adapt to the requirements of the current project.

2.3.1 Purpose

The tactical layer is modelled in order to analyse the delay propagation between flights and adaptability from disruptions (cancellations, background and foreground delay) in the ATM environment, limited resources (fixed number of slots). Delay propagation depends on the capacity of the system, either the airspace capacity, airport capacity or airline capacity. Airport capacity is explicitly modelled through airport departure and arrival slots modelling, the others are implicitly modelled. On the other hand, adaptability from disruptions is a fundamental attribute of the ATM environment in order to be a resilient system. Both features are evaluated using stochastic models that change the behaviour of the system each simulation run.

2.3.2 Model details

2.3.2.1 Flight

Flights are modelled in the tactical layer as data structures that merges flight plan data with references to the passenger itineraries related to that flight plan. Flight plans include flight schedule data, such as scheduled times and references to the stakeholders involved (airline, airports). Flight plans also include different flight times in order to study the model uncertainties in taxi times, runway times and en-route times. Finally, several delays are considered in order to analyse reactionary delay, departure delay and arrival delay.

Not every flight is an object of study. Flights are flagged as excluded if we find they are difficult to extract conclusions from its state, for example, dealing with circular flights and military flights, but they are included in the simulation because they are important to model airport capacity.

Table 1 describes the state variables of the model, as well as some of their attributes.

Table 1. Tactical model, flight state variables

State variable	Cardinality	Description
Identifier	1	Flight identifier (UNIQUE)
ADEP	1	Departure airport ICAO code
ADES	1	Arrival airport ICAO code
Aircraft	1	Aircraft type code
Airline	1	Airline ICAO code
SOBT	1	Scheduled Off block time
SIBT	1	Scheduled In block time
state	1	Current simulation status of the plan
wait4aircraft	1	Reactionary delay
wait4pax	1	Wait for pax time

State variable	Cardinality	Description
wait4offblock	1	Wait for non-reactionary delay
taxi_out_time	1	Taxi out time
wait4depslot	1	Wait for departure slot time
runway_out_time	1	Runway out time
flight_duration	1	Flight duration
en_route_delay	1	En-route variability
en_route_recover	1	Actual departure delay recovered en-route
wait4arrslot	1	Wait for arrival slot time
runway_in_time	1	Runway in time
taxi_in_time	1	Taxi in time
Passengers	-	Passenger list
Max seats	1	Number of seats to be fully loaded
excluded	1	To exclude the flight from metrics (difficult to make assumptions)
cancelled	1	Cancelled flights (flight execution times not available)

The initial flight plan state expects the flight will arrive on time (zero delay). Flight plan is updated accordingly by the simulation of different flight events.

Some derived times are simulated in simulation time as follows in (Table 2).

Table 2. Tactical model, flight derived simulation times

Signature	Description	Pseudocode
ART	Get aircraft ready time	SOBT + wait4aircraft + wait4pax
EOBT	Get estimated off-block time	ART + wait4offblock
ETOT	Get estimated take-off time	EOBT + taxi_out_time + wait4depslot + runway_out_time
ETA	Get estimated time of arrival	ETOT + flight_duration + en_route_delay + en_route_recover + wait4arrslot
EIBT	Get estimated in-block time	ETA + runway_in_time + taxi_in_time
AOBT	Get actual off-block time	EOBT if status > AIRCRAFT READY

Signature	Description	Pseudocode
ATOT	Get actual take-off time	ETOT if status > TAXI_OUT
ATA	Get actual time of arrival	ETA if status > EN ROUTE
AIBT	Get actual in-block time	EIBT if status > TAXI_IN

2.3.2.2 Passenger itineraries

Passenger itineraries are modelled in the tactical layer as data structures that merges itinerary data with references to the flight plans where the itinerary takes place. Different implementation mechanisms ensure the double link between those elements is not broken.

A passenger itinerary includes the number of passengers that follows the same itinerary, as well as some ticket attributes as the fare type and fare cost. Consequently, there might be more than one instance of the same itinerary depending on the number of ticket types the itinerary is offered. Other scheduled itinerary details such as the original schedule and original route is referenced for analytical purposes.

Table 3 describes the state variables of the model, as well as some of their attributes.

Table 3. Tactical model, passenger itineraries variables

State variable	Cardinality	Description
code	1	Generated passenger code
size	1	Number of passengers in the group
ticket type	1	Ticket type (premium, standard)
itinerary	N	List of itinerary flights
origin	1	Origin airport
destination	1	Destination airport
sched_departure	1	Scheduled departure time for first leg
sched_arrival	1	Scheduled arrival time for last leg
reallocated	1	Group are reallocated or not

The list of flight plan references that describe the passenger itinerary is the main state variable that derives other values, such as passenger delay, times and categorisation. These references make possible the estimation of passenger delays and times during the simulation to adapt the behaviour of other models. Not only are these references used, but the ticket state variable and the scheduled itinerary ones are used in combination with the list of flight plan references to categorise the passengers for calculating reactionary delay and to estimate the time a passenger will stay at the airport gate to catch a later flight.

Table 4. Tactical model, passenger derived simulation times

Signature	Description	Pseudocode
is_onward_long_haul	Can be categorised as an onward long haul itinerary	From current leg onwards, check if there is a long haul leg
is_premium_conn_pax	Can be categorised as a premium connecting pax	Check if the fare type is premium and the itinerary has more than 1 leg
departure_delay	Get departure delay at origin	AOBT first leg - sched_departure
arrival_delay	Get arrival delay at destination	AIBT last leg - sched arrival
estimated_pax_at_gate_time	Get time to be ready for the next leg (estimated)	EIBT prev leg + connecting time
actual_pax_at_gate_time	Get time to be ready for the next leg (estimated)	AIBT prev leg + connecting time

2.3.2.3 Event overview and scheduling

The tactical layer runs multiple discrete event simulations. A discrete event simulation explicitly specifies the times where the model state can be updated to improve the simulation performance. The implementation ensures that these discrete events are triggered in the correct order by using a monotonic simulation clock and a simulation manager, managing the scheduling and dispatching of events. The implementation of both components is also included in the library.

The event chain diagram below (Figure 13) identifies the list of discrete events considered in the tactical model as well as the dependences between them.

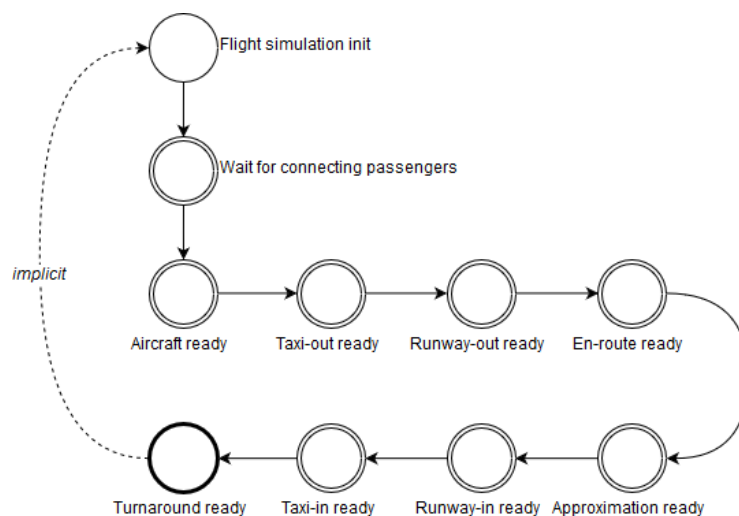


Figure 13. Tactical layer event chain diagram

A dedicated flight simulation ('init' event) is created at the same time a flight instance is loaded with input data. The implicit link between the turnaround ready event and the flight simulation init event ensures the next flight using the same aircraft actually waits for the former event to be succeeded.

(i) Flight simulation init event

The flight simulation init event performs several auxiliary task in order to propagate delay from previous flights.

1. First, it subscribes to the aircraft ready event that is triggered from the previous flight if that flight exists and waits until scheduled departure time. The last event that occur resumes the simulation of the flight.
2. Flight reactionary delay is updated as soon as the simulation is resumed.
3. The event checks if the flight is cancelled.
 - a. In case the flight is cancelled, the airport connectivity graph is updated, all passengers are passed to the reallocation model and the next flight is notified in order not to wait for it.
 - b. In case the flight is not cancelled, a Waiting for connecting passengers' event is scheduled.

(ii) Waiting for connecting passengers' event

This event calls the wait for passengers model to set how much reactionary delay should be assumed because of connecting passengers. It is an iterative process:

1. Calculates how much time will wait for those passengers.
2. Explicitly waits for that time through the event scheduler.
3. Flight reactionary delay is updated as soon as the simulation is resumed.
4. The simulator checks if those passengers have actually arrived or not and propose a new time wait.
 - a. If wait time is positive and the accumulative delay is less than max_wait time , returns to step 2.
 - b. If wait time is zero, passengers that miss the connection are rescheduled to other flights, airport connectivity graph is updated and then an aircraft ready event is scheduled.

(iii) Aircraft ready event

As soon as the aircraft is ready, the simulator calls the non-reactionary delay model to calculate the amount of delay that is caused by an ATFM regulation or other disruptions.

1. Calculates how much time is required to wait before taxi-out takes place.
2. Explicitly waits for that time through the event scheduler.
3. Flight non-reactionary delay is updated as soon as the simulation is resumed.

4. A Taxi ready event is scheduled.

(iv) Taxi-out ready event

As soon as the aircraft is ready for taxi-out, the simulator calls the taxi model to calculate the amount of time required to perform taxi operations in the airport of origin.

1. Calculates how much time is required to taxi-out.
2. Explicitly waits for that time through the event scheduler.
3. Taxi out estimation is updated with actual time.
4. A runway-out ready event is scheduled.

(v) Runway-out ready event

As soon as the aircraft is ready for runway-out, the simulation calls the airport capacity model to wait for a take-off slot and occupy the runway.

1. Request a take-off slot.
2. Explicitly waits until the runway is ready for the flight to be used.
3. Departure management delay is updated as soon as the simulation is resumed.
4. Explicitly waits until the occupation time is elapsed through the event scheduler.
5. Runway out time estimation is updated with actual time.
6. The take-off slot is released to be used for other flights.
7. An en-route ready event is scheduled.

(vi) En-route ready event

As soon as the aircraft is ready for en-route, the simulation calls the en-route time and delay recovery model to calculate actual en-route time.

1. Calculates how much time is required for en-route.
2. Explicitly waits for that time through the event scheduler.
3. Flight time estimation is updated with actual time.
4. An approximation ready event is scheduled.

(vii) Approximation ready event

As soon as the aircraft is ready for approximation, the simulation calls the airport capacity model again to wait for a landing slot in the air and then occupy the runway.

1. Request a landing slot.
2. Explicitly waits in the air until the runway is ready for the flight to be used.
3. Arrival management delay is updated as soon as the simulation is resumed.
4. Explicitly waits until the occupation time is elapsed through the event scheduler.

5. Runway in time estimation is updated with actual time.
6. The landing slot is released to be used for other flights.
7. A taxi-in ready event is scheduled.

(viii) Taxi-in ready event

As soon as the aircraft is ready for taxi-in, the simulator calls the taxi model again to calculate the amount of time required to perform taxi operations in the airport of destination.

1. Calculates how much time is required to taxi-in.
2. Explicitly waits for that time through the event scheduler.
3. Taxi in estimation is updated with actual time.
4. A turnaround ready event is scheduled.

(ix) Turnaround ready event

As soon as the aircraft is ready for turnaround, the simulation calls the Reactionary delay model to calculate the time required for the turnaround process and to notify the next flight.

1. Calculates how much time is required to turnaround.
2. Explicitly waits for that time through the event scheduler.
3. Turnaround time estimation is updated with actual time.
4. The next flight is notified.

2.3.2.4 Behavioural details

(i) Reactionary delay

Reactionary delay is modelled explicitly by propagating the delay of previous flights but also modelling the rules followed by the airlines in order to decide if a connecting passenger is waited or not.

Delay propagation is modelled using event subscription, one of the features of the model framework.

- As described above, the simulation is driven by an event scheduler.
- Each time an event is performed, there is a list of triggers that are enabled in order to notify other process that the event was successfully simulated.
- Each event has to be subscribed explicitly in order to be notified. This subscription is done when the model is created.
- Delay propagation depends on turnaround times. Turnaround times are assigned to each flight as an additional status variable in the flight model.

Wait for connecting passengers delay is modelled using a rule-based model (see also Table 5) based on:

- Do not wait unless the flight load value is less than certain threshold, defined as an input of the model.
- Delay is bounded up to a certain number of minutes defined as an input of the model that should be calibrated.
- Depending on the factors, not every passenger is waited:
 - Only premium ticketed connecting pax.
 - Only onward long-haul passenger.

Table 5. Tactical model, passenger wait variables

State variable	Cardinality	Description
max wait	1	Max accumulated waiting time for connecting passengers
load_threshold	1	Percentage of seats occupied

Pax waiting time depends on the previous flight arrival but also the passengers connecting time. This value depends on the airport and two options are considered:

- Connecting time in the same airport terminal.
- Connecting time between different airport terminals (for international flights).

Table 6 describes the attributes of the model where this information is collected.

Table 6. Tactical model, passenger connecting time attributes

Field	Type	Description
airport id	Key	Airport ICAO id
mct_dom	Column	Minimum connection in the same airport terminal
mct_intl	Column	Minimum connection time between different airport terminals

(ii) Non-reactionary delay

The non-reactionary delay depends on ATFM regulation but also in some events in the environment that are difficult to model. This model reproduces both events using a single stochastic model. This model will be calibrated for the whole dataset.

(iii) Passenger reallocation

Passenger reallocation is modelled explicitly using connectivity graphs and reallocation rules. The connectivity graph is used to check if there is a flight route between two airports, modelled as a table of tables.

Table 7. (Tactical) Parent airports table

Field	Type
Airport id	Key
Departures table	Column
Arrivals table	Column

Table 8. (Tactical) Nested airports table

Field	Type
Flight code	Key
Origin	Column
Destination	Column

There is a method to find existing routes between airports that performs a full search of the two affected tables (origin departures and destination arrivals) and looks for flights with same flight code that follows the same route as requested (direct flights) or two flight codes that shares the same intermediate airport that follows the same route as requested (one-stop flights). There are also two additional methods to update the connectivity graph in order not to provide cancelled routes or already arrived routes. Reallocation rules depend on the type of airline and airline alliances as well as the passenger fare type. In particular:

- Non FSC standard passengers are rescheduled to any other flight from the same alliance.
- FSC standard passengers are rescheduled to any other flight if the passenger has also an arrival delay greater than 5 hours.
- Other passengers (FSC premium) are rescheduled to any other flight with no restrictions.

The reallocation model holds a table with airline data, such as airline type and alliances between airlines.

Table 9 contains the fields of the airlines.

Table 9. Tactical model, passenger connecting time attributes

Field	Type	Description
Airline	Key	Airline ICAO code
Type	Column	Airline type (REG, FSC, LCC, CHT)
Alliance	Column	Alliance name
Partner	Column	Partner name

The model provides a method that given a list of passengers, they are reschedule on other flights and those that are not eligible for reschedule are returned. Those passengers are stored in a queue to be reschedule as last-minute passengers if there are some passengers that miss the flight. There is one queue per airport as described in Table 10. The queue is updated when a reallocation process takes place.

Table 10. Tactical model, airports waiting passengers table

Field	Type
Airport id	Key
Waiting passenger codes (list)	Column

There is a method to update the queue when a reallocation process takes place.

(iv) Taxi time

This stochastic model is described through a Normal probability distribution function calibrated with input data. Taxi in times and taxi out times are calibrated separately and depends on the airport. The fitting parameters are provided as inputs of the model using a table such as the one in Table 11.

Table 11. Tactical model, taxi time table

Field	Type
Airport id	Key
Taxi out mean	Column
Taxi out standard deviation	Column
Taxi in mean	Column
Taxi in standard deviation	Column

(v) Airport capacity

This model considers two different aspects: there is a slot where the runway is exclusive used for a single flight and flights have to wait until the runway is ready for them. The main attribute that defines a runway slot is the occupation time. There is a single occupation time for both operations

(departures and arrivals) per airport. It is considered as an input of the model. The model framework provides an implementation of a mutual exclusive use of the runway. This auxiliary model includes a request queue to have several flights waiting at the same time. The take-off event has to subscribe explicitly to the event to be notified when the runway is ready to be used, as described in Table 12.

Table 12. Tactical model, airport runway table

Field	Type
Airport id	Key
Occupation time	Column
Runway model	Column

(vi) En-route time and delay recovery

En-route time is modelled stochastically describing a Normal probability distribution. The distribution is centred at the en-route time. This en-route time is considered as an input of the model but also the variability of the value. Delay recovery is explicitly model as a function with two variables:

- amount of residual delay not to be recovered.
- maximum time to be recovered defined as a percentage of the flight time.

Both values have to be calibrated.

(vii) Flight cancellations

The cancellation of flights is modelled as a stochastic event with a probability ratio to be affected and modelled as a Uniform distribution. The probability ratio is considered as an input of the model. There is a single model for the whole simulation.

3 Initial 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 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.

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 the ref. [1] 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_0 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 [2] 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 [3] 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 [4]), we used a linear rule of thumb to compute their respective proportions:

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

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:

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

(iv) Airline

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 we were directly fed the Flight agent.

From BADA model, we also extract an average MTOW of each OD pair (in fact, an average $\sqrt{\text{MTOW}}$). 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 same high-level value for all airlines, taken from [5], 108 euros/minute/flight.

(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

To calibrate the ANSPs, we used the assumptions that the ANSPs were almost 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 used the average cost per flight hours controlled from PRU reporting [6] to infer their capacity by equating their costs and their revenues. More specifically, for a given ANSP, we use their unit rate u and the historical traffic to compute their revenues r , that we assume is equal to their cost c . The delivered capacity is thus c/ω , where ω is the cost per unit of capacity. Given the initial delay δt_0 computed also from historical data, we can then fix the parameter A as following:

$$A = \delta t_0 \exp\left(-\frac{T}{C}\right)$$

where T is the initial historical traffic level. 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

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 analysis, where one changes a parameter in the model (usually not calibrated on data) to see how it affects the model.

In the model we have very few free behavioural parameters, most of them being already inferred from data one way or another. This phase requires heavy simulation runs, as one needs to sweep the parameters, compute some metrics, compare them to a metric from the data, etc. Moreover, the model needs to be quite stable to perform such analysis, otherwise the tuned parameters are obsolete when a part of it is changed. For these reasons, we are only now ready to perform such an analysis, and it will be done before the next deliverable D5.2. We plan to perform the following tests, tuning the:

- competition parameter for flight plans (intensity of choice based on cost of flight plan);
- competition parameter for demand, to choose among itineraries;
- rate of change of price;
- increase of demand, not only using income (in particular to match STATFOR predictions, see Section 5.1.1).

One of the important things to test is the stability of the model with respect to these parameters. As mentioned several times, the model can indeed be unstable, e.g. losing cohesion in the sense that metrics are grow infinitely or oscillate, or even going into a chaotic mode. This can often happen with this kind of dynamical model, in particular when lagged variables are present.

Other checks can be performed too in this phase, in order to cross-check some facts and see if there is an issue somewhere in the model. We have started doing this, and in Figure 14 we show for instance 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 [12], and on the right hand side we show a similar plot obtained with the model, using the ‘current’ scenario.

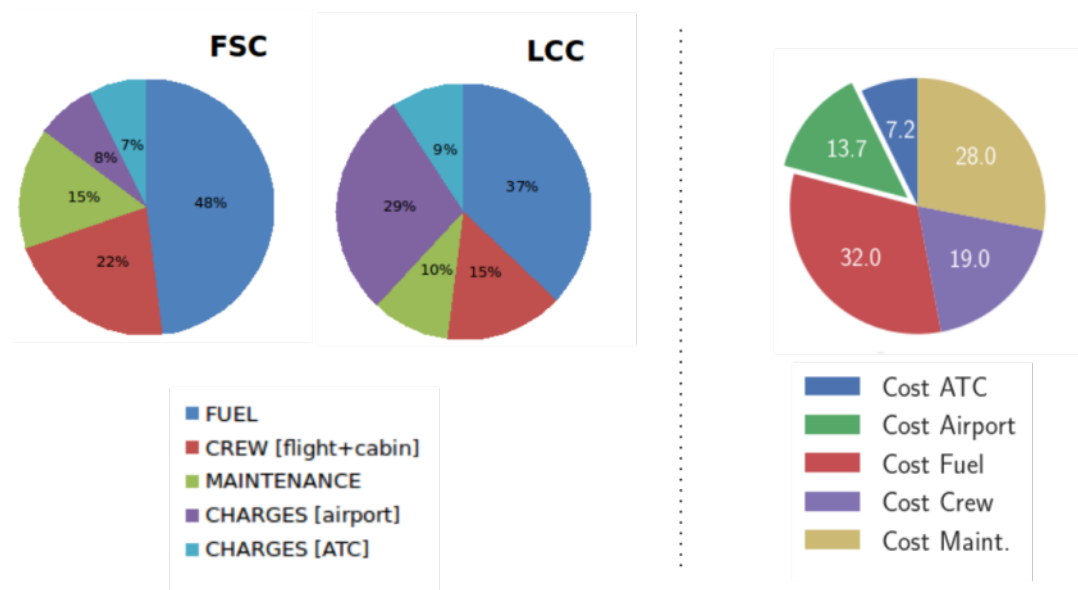


Figure 14. Airline cost data comparison between empirical data and outputs of the model
Empirical data source: ICAO DATA+ “air carrier finances” [12]

Keeping in mind that the right-hand side is an average over all the types of airlines, the agreement between both is quite good. Some differences still exist, for instance the cost of fuel seems to be underestimated in the model with respect to reality. 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 cost seems to be overestimated in the model. More checks and further refinements will be performed in Deliverable 6.3 and 5.2, as flagged in the list of features in Section 6.2.

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 is 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.

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 15. 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 15. 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 16 shows how many of these elementary taxons can be found in the data.




Figure 16. Distribution of elementary taxes, characterised by number of airports (abscissa)

The vast majority of these taxes 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 taxes 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. As explained before, these 3 elementary taxons can then generate a high number of taxes by combining them. It is thus important to see what the most represented taxes are in the data, which is represented in Figure 17.




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

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 forth). 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). For the schedule mapper, we selected all the taxons represented above. These taxons, together with their distribution, are recorded and passed to the schedule mapper.

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.

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

Figure 18 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 19 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 19. Distribution of turnaround times (in hours) for taxons (1, 0, 0), (2, 0, 0), and (7, 0, 0)

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 20 a) presents all the historical routes between EBBR and LTBA (in this case there are 26 flights in the period analysed). Figure 20 c) presents the result of the application of the clustering reducing the number of routes to 12. As presented in Figure 20 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.

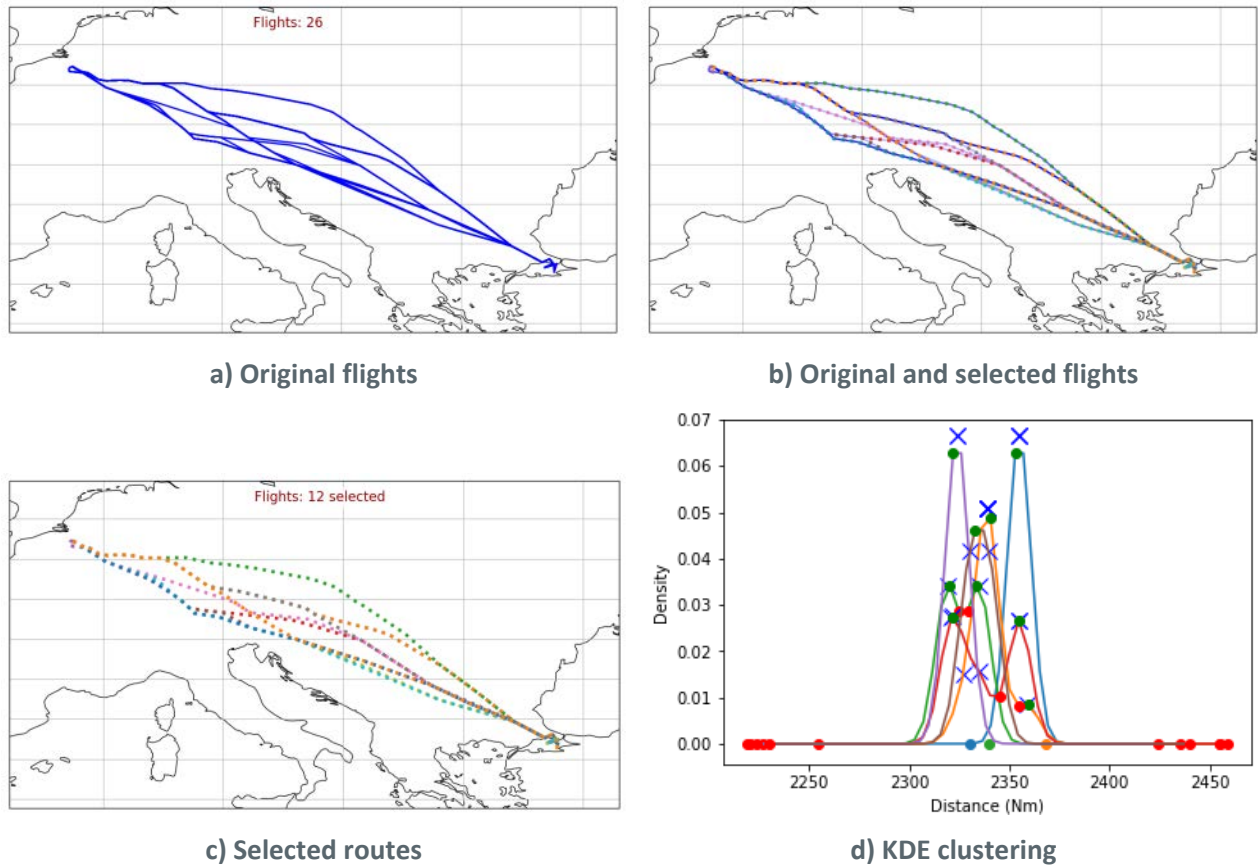


Figure 20. Example of route clustering EBBR – LTBA

Figure 21 presents another example between DGAA and EGLL. In this case, it is easier to observe in Figure 21 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 a 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.

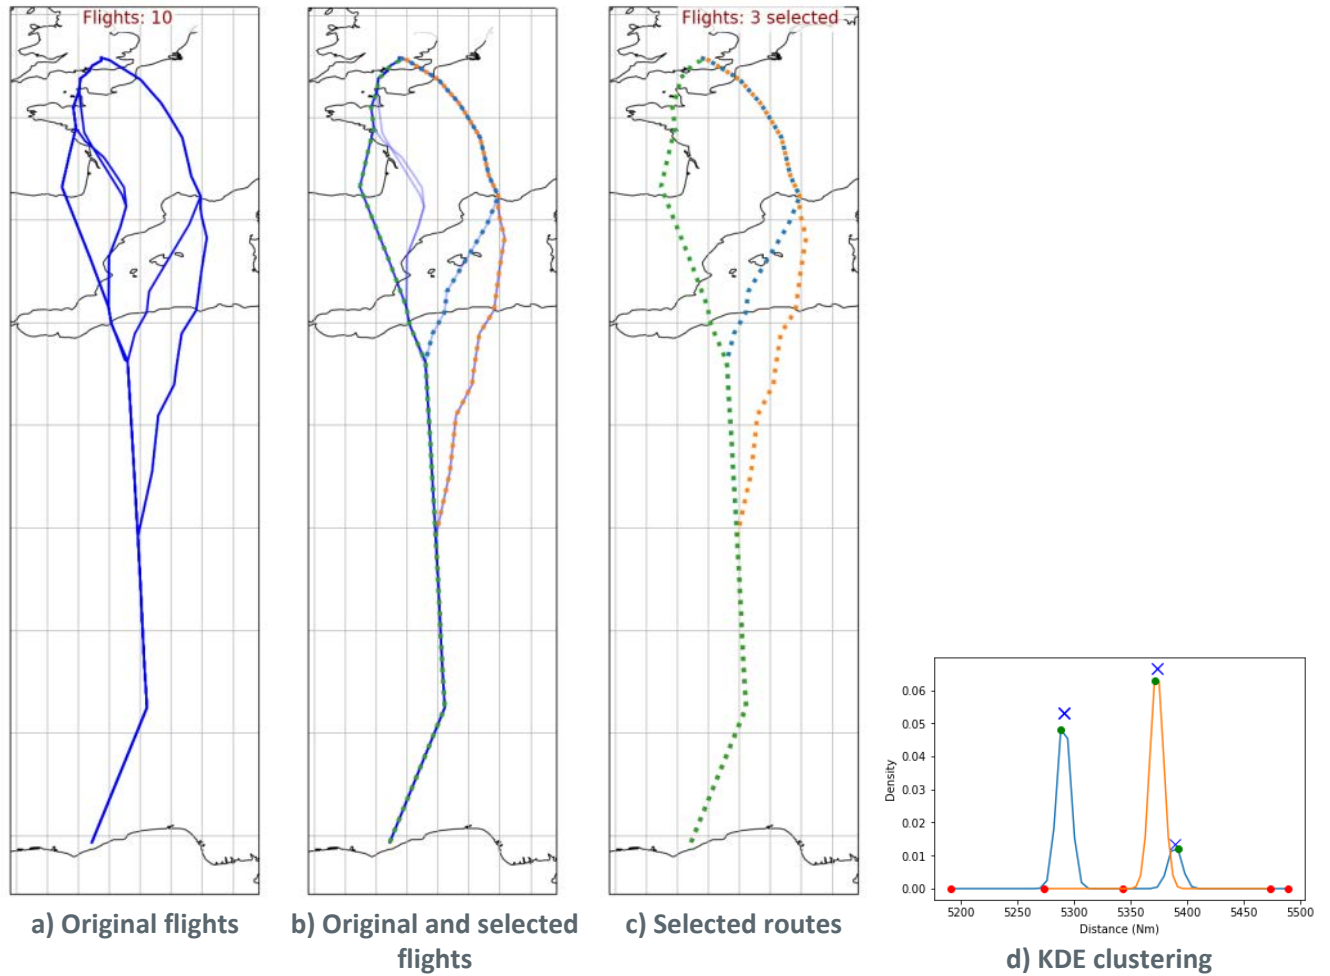


Figure 21. Example of route clustering DGAA – EGLL

(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 22 shows two examples of trajectories with their associated speed and FL requests.




Figure 22. 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 23 shows the average flight level computed for the examples of Figure 22.




Figure 23. 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 24 for the B737.




Figure 24. 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 25 for the A320.

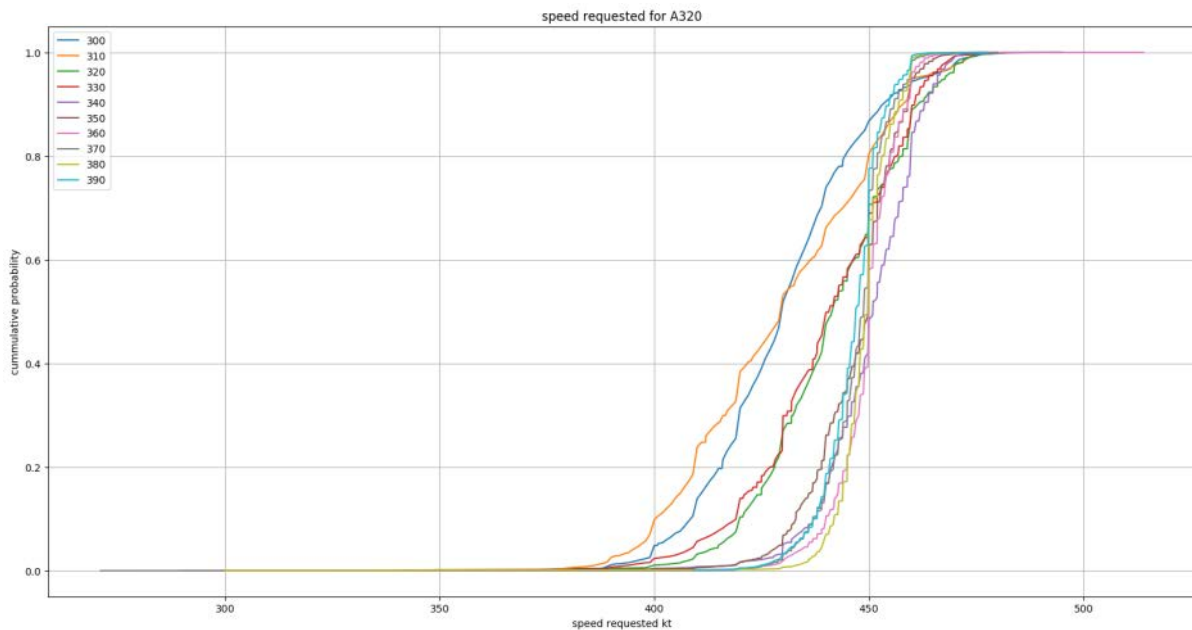


Figure 25. Example of distributions of speed requested for A320 by flight level

(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.






Figure 26. 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 26 shows the climb and descent estimated for the two example flights.



a) Climb distance



b) Descent distance

Figure 27. 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 27 shows these probabilities for the climb and descent distances for the B737.

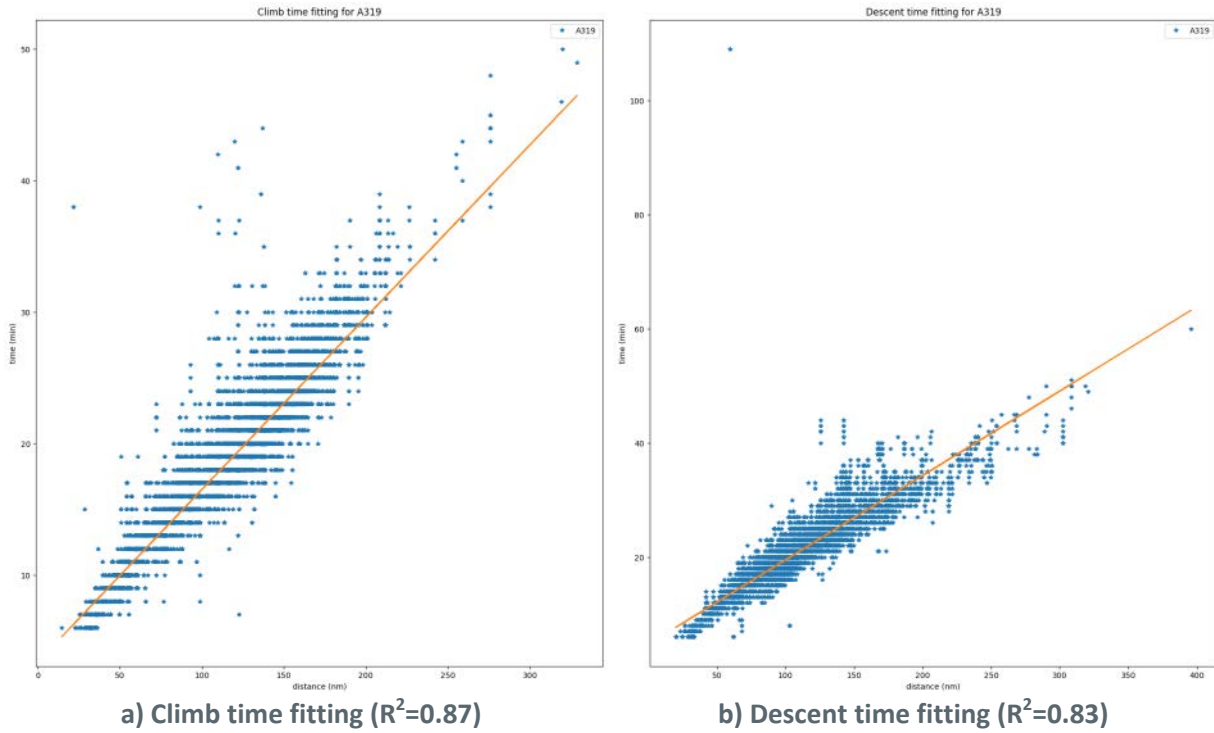



Figure 28. Example of linear fitting of climb and times for A319

For each aircraft type a linear fitting between the climb and descent distance and the time required to perform them is computed. Figure 28 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.



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 29, 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 29 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 [11].

(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 30 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 31).




Figure 30. Flight plan distances: historical flights and average for FP selected per schedule




Figure 31. Histogram of flight plan distance for all FPs generated




Figure 32. Cumulative probabilities of FP distance (NM)

Figure 32 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 13 shows some statistics on the flight plan distance for the historical flights, all the generated and the selected.

Table 13. 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 33. Taxi times and schedule buffers for historical and selected FPs




Figure 34. Cumulative probability of taxi and buffer time

Table 14. 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 33 the times generated for the flight plan selected are very similar to the historical ones. This is further observed in Figure 34 and in Table 14, 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 34.





Figure 35. 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 than 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 not 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 36. Fuel as a function of FP distance (NM) per aircraft type


Figure 36 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 37 where for long flights (greater than 3000 NM) there is a lower cost on charges than for shorter ones. See in Figure 37.b and Figure 37.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 37. 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 38 shows the revenue on en-route airspace charges per NAS.




Figure 38. 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 [9]. 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 39. Figure 40 shows the difference in service units which consider aircraft distance and size within the different NAS.




Figure 39. En-route revenues per NAS from FPs generated and reported service units




Figure 40. 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

From all the days in the period AIRAC1313-AIRAC1413, those where industrial actions were produced are filtered out to avoid an unrealistic ATFM delay distribution of delay. For all the remaining traffic, the probability of going through a regulation and that regulation being the most penalising for the flight is estimated giving a probability of 5.6% of being affected by an ATFM regulation.




Figure 41. Probability distribution of ATFM delay

Then all the ATFM delay is used to generate a cumulative probability distribution of the amount of delay assigned as presented in Figure 41. Note that even if the flight goes through a regulation there is still 40% probabilities of having a delay of zero minutes assigned.

3.2.2.3 Post-calibration of ATFM generator

The post-calibration of the ATFM Regulation Generator is dependent on the actual delay that is assigned to flights. As this block produces the probabilities of delay to be used by the tactical layer, a feedback loop is required. The ATFM regulator generator needs the outcome of the tactical execution of the traffic with the regulations to assess if the probabilities are adequate comparing the results to historical delay data. This has started to be done with the first runs of the tactical layer and will continue in the next stages of the development of the model.

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

The main data sources that will be used for the calibration of the Flight Plan Generator are passenger itineraries for 12SEP14, include IATA Passenger Intelligence Services (PaxIS) and GDS data for the passenger information, and ACI EUROPE and airline data for the airport connectivity and flights load factors.

3.3 Tactical layer

The calibration of the tactical model was performed using several data sources from previous work (in SESAR1, ComplexityCosts). Those data sources were transformed to generate:

- Flight schedule data, such as origin, destination, airline, flight identifier, departure and arrival schedule times, aircraft type, maximum seats and registration number;
- Taxi times;
- Airline classifications (per type and alliances);
- Airport classifications;
- Connection times;
- Turnaround times;
- Airport capacity values;
- Wake turbulence aircraft type categorisation;
- Passenger itineraries.

The calibration of the model considers the data from the 12SEP14 as the baseline scenario. As described on the tactical layer, these are the model components susceptible to be calibrated:

- Airline model;
- Airport model, that includes:
 - Taxi model;
 - Runway model;
 - Flight connection model;
- Event component;
- Flight schedule model.

There are some others, such as the passenger itinerary model, that are simple data placeholders, so they are not calibrated, just loaded with data.

3.3.1 Loading required input data to perform a simulation of all the events

Passenger itineraries for the day used for calibration by the tactical layer are based on the analysis done in ComplexityCosts. This previous SESAR1 project processed anonymised passenger itineraries (derived from various sources including GDS passenger data and IATA Passenger Intelligence Services (PaxIS) data), calibrated with ACI EUROPE and airline data for airport connectivity and flight load factors. The passengers are grouped by passengers following the same itinerary, fare type and fare cost. The simulation model also uses passenger groups to simulate scenarios more efficiently, so there is a direct correspondence between passenger group model and input passenger itinerary data (as shown in Table 15).

Table 15. Passenger itinerary data and group model

Passenger itinerary data	Passenger group model
code	code
pax	size
premium/standard	fare_type
legN	itineraries (list)

For the calibration of the tactical layer, historical flight schedules obtained from Innovata have been analysed and mapped to the flights of the day tested (12SEP14). The mapping is as described in Table 16.

Table 16. Flight schedule mapping

Flight schedule data	Flight model
AOnew	airline
ADEP	origin
ADES	destination
max_seats	max_seats
FltNum	code
SOBT.UTC	SOBT
SIBT.UTC	SIBT

Once both models are loaded with data, the simulation of all events can be carried.

3.3.2 Calibrating the off-block delay model

3.3.2.1 Reactionary delay

- The model considers reactionary delay as the delay generated because the required aircraft is not ready but also the delay generated to wait passengers from a connecting flight that is delayed.
- Aircraft ready time is modelled using event subscription. Aircraft ready time is not set until the turnaround event from the last flight of the previous flight is succeeded.
- In order to set flight interdependence because of using the same aircraft, flight schedule data are grouped by Registration number and then sorted by scheduled departure time. The sorting also checks that the destination airport of one flight and the origin airport of the other matches.
- Turnaround times were estimated from ALL_FT+ data by computing the turnaround time distributions using tail numbers, finding that the major features determining the turnaround times are: wake turbulence category, size of the airport and airline type.
- Uncertainty in actual turnaround times is modelled as an Exponential distribution defined by $\lambda = 15$ minutes (estimated from ALL_FT+ dataset). There is a random sample generated for each flight.
- Reactionary delay also considers waiting up to 10 minutes until the flight has 80% of expected passengers.
- Connecting times between terminals are estimated from ALL_FT+ buffers between flight schedules.

3.3.2.2 Non-reactionary delay

- The model considers non-reactionary delay as the delay generated because of an ATFM regulation or any other reason that prevents the flight from leaving the gate.
- ATFM regulation delay depends on an empirical cumulative distribution function and a probability of having an ATFM regulation. Both values were calibrated by historical schedules with ATFM regulation data removing regulations applied for industrial actions. The parameters are:
 - $p = 0.056$
 - Non ATFM delay is modelled as an exponential distribution defined by $\lambda = 15$, truncated to 90 minutes (estimated from DDR dataset). The probability of having non ATFM delay is 0.15.

3.3.3 Passenger reallocation model

- The airport connectivity graph is extracted from input flight schedules.
- Airline classification data are extracted from public sources.

3.3.4 Simulated taxi times

Taxi times modelled as a normal distribution for each airport considered in the model. Mean and standard deviation times were extracted from the CODA dataset for the same year.

3.3.5 Limiting airport capacity through a data driven approach

- Airport capacity value extracted from input flight data, assuming 1 runway per airport and considering departures and arrivals as a whole.
- Operations are then grouped into 1-hour slots overlapped each 15 minutes.
- Airport capacity is later statistically inferred using the 98.2 percentile of the number of operations per airport slot considering all slots of an airport.
- The occupation time per airport is finally considered as the inverse of capacity.
- This value is pre-calculated and included as an input of the runway model.

3.3.6 En-route time

- Flight duration is approximated by the last filled flight plan in ALL_FT+
- Recovery rule of thumb (based on AO consultation and consortium expertise):
 - Recover delay to a residual of 5 minutes
 - Up to 5% of flight duration
- Flight duration: Normal distribution of mean centre to flight duration and 0.02 of standard deviation.

3.3.7 Flights affected by cancellations

Operational cancellations in 2014 remained stable at 1.5%, with peaks of around 8% on days with industrial action – as shown in Figure 42, these days correspond to airline strikes [7]. Figure 42 shows the cancellation rate experienced on a daily basis, industrial actions are marked at some peaks. In MAY14, French ATC industrial action generated peaks at 2% of cancellations and the industrial actions of 24-26JUN14 increased the cancellation rate up to 5% [7].




Figure 42. Average daily cancellations 2014

Source: CODA Digest [7]

Figure 43 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 43. Monthly share of operational cancellations 2014

Source: CODA Digest [7]

4 Scenarios and factor modelling

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 tech. development in 2035), 'H35' (high economic growth and high tech. development in 2035), 'L50', and 'H50' (*idem.* in 2050). On top of that, we selected four foreground factors to be modelled for this deliverable:

- BEO1, modifying the price of fuel with 5 possible values: 0.1 (LL), 0.3 (L), 0.5 (D, current), 1.0 (H) and 2.0 (HH), in euros per kilogram of fuel. Currently, in the tactical layer, the actual fuel price is not directly used as input but rather its effect on daily operations. There are two scenarios considered a low fuel cost scenario marked as 'A' and a high fuel cost named as 'B'. In the final model, after calibration, both scenarios will be referred to 'LL', 'L', 'D', 'H' or 'HH' according to the numeric fuel prices.
- BEO2, modifying the CRCO charges and the airspace structure, with 6 possible values: 'D' (current), 'Homo' (homogeneous charges in the ECAC space), 'FAB' (homogeneous charges per FAB), 'Zero' (unit rates set to zero for ECAC space), 'Double' (unit rates doubled for ECAC space), 'Imbalance' (unit rates double for ECAC spaces which have a unit rate above the average and unit rates halved for ECAC spaces which have a unit rate below average). Note that the scenarios where the unit rate is changed are executed as test for the Flight Plan Generator block as in the whole integrated system, those values will be the outcome of the Economic model.
- BTS5, the implementation level of 4D trajectories, with 4 possible values: 'D' (current), 'L', 'M', and 'H' (see next for a reminder of how the SESAR technological improvements impact the models). Additionally the tactical layer considers two scenarios for BTS5, scenario 'A' denotes partial implementation while 'B' represents complete deployment according to SESAR performance targets.
- BTO4, the implementation level of passenger reaccommodation tools, with parameters: 'A' basic passenger tools, and 'B' advanced passenger management tools (see Section 4.5 for more details on how it has been implemented.)

Table 17 shows all the scenarios (56) which have been run so far by at least one block.

Table 17. Scenarios simulated by Vista

The last column includes the blocks which have simulated this scenario already (EM is economic model, FPG is flight plan generator, MER is Mercury).

Scenario	Description	Run by
Current	Current situation (SEP14)	EM FPG
Unitary_BEO1_L_Current	Low fuel price in current env.	FPG
Unitary_BEO1_H_Current	High fuel price in current env.	FPG
Unitary_BEO1_LL_Current	Very low fuel price in current env.	FPG
Unitary_BEO1_HH_Current	Very high fuel price in current env.	FPG
Unitary_BEO2_FAB_Current	FAB structure in current env.	FPG
Unitary_BEO2_Homo_Current	Homogeneous ATCO charges in current env.	FPG
Unitary_BEO2_Zero_Current	Unit rate of EUROCONTROLS' CRCO members state set to zero	FPG
Unitary_BEO2_Double_Current	Unit rate of EUROCONTROL's CRCO members state set to double in value	FPG
Unitary_BEO2_Imbalance_Current	Unit rate of EUROCONTROL's CRCO members state which are higher than average unit rate doubled Unit rate of EUROCONTROL's CRCO members state which are lower than average unit rate halved	FPG
Unitary_BEO1_A_Current	Low fuel prices affecting recovery strategy	MER
Unitary_BEO1_B_Current	High fuel prices affecting recovery strategy	MER
Unitary_BT55_A_Current	Basic 4D trajectories implementation	MER
Unitary_BT55_B_Current	Advanced 4D trajectories implantation	MER
Unitary_BTO4_A_Current	Better passer reaccommodation tools	MER

Current

Scenario	Description	Run by
Unitary_BTO4_B_Current	Advanced passenger reaccommodation tools	MER
Dual_BEO1_L_BEO2_FAB_Current	Low fuel price & FAB structure in current env.	FPG
Dual_BEO1_L_BEO2_Homo_Current	Low fuel price & Homogeneous ATCO charges in current env.	FPG
Dual_BEO1_H_BEO2_FAB_Current	High fuel price & FAB structure in current env.	FPG
Dual_BEO1_H_BEO2_Homo_Current	High fuel price & Homogeneous ATCO charges in current env.	FPG
Dual_BEO1_BTS5_A_Current	Low fuel prices affecting recovery strategy and basic 4D trajectories	MER
Dual_BTS5_BTO4_A_Current	Basic 4D trajectories and better passenger reaccommodation tools	MER
Dual_BEO1_BTO4_A_Current	Low fuel prices affecting recovery strategy and better passenger reaccommodation tools	MER
Dual_BEO1_BTS5_B_Current	High fuel prices affecting recovery strategy and advanced 4D trajectories	MER
Dual_BTS5_BTO4_B_Current	Advanced 4D trajectories and passenger reaccommodation tools	MER
Dual_BEO1_BTO4_B_Current	High fuel prices affecting recovery strategy and advanced passenger reaccommodation tools	MER
Multi_BEO1_BTS5_BTO4_A	Low fuel prices affecting recovery strategy, basic 4D trajectories and better passenger reaccommodation tools	MER
Multi_BEO1_BTS5_BTO4_B	High fuel prices affecting recovery strategy, advanced 4D trajectories and advanced passenger reaccommodation tools	MER
L35_baseline	Baseline environment in 2035 with slow eco. growth and slow tech. development.	EM

2035

Scenario	Description	Run by
Unitary_BEO1_LL_L35	Very low fuel price in low 2035 baseline	EM
Unitary_BEO1_L_L35	Low fuel price in low 2035 baseline	EM
Unitary_BEO1_H_L35	High fuel price in low 2035 baseline	EM
Unitary_BEO1_HH_L35	Very high fuel price in low 2035 baseline	EM
Unitary_BT55_L_L35	Low implementation of 4DT in low 2035 baseline	EM
Unitary_BT55_H_L35	High implementation of 4DT in low 2035 baseline	EM
H35_baseline	Baseline environment in 2035 with fast eco. growth and fast tech. development.	EM
Unitary_BEO1_LL_H35	Very low fuel price in high 2035 baseline	EM
Unitary_BEO1_L_H35	Low fuel price in high 2035 baseline	EM
Unitary_BEO1_H_H35	High fuel price in high 2035 baseline	EM
Unitary_BEO1_HH_H35	Very high fuel price in high 2035 baseline	EM
Unitary_BT55_L_H35	Low implementation of 4DT in high 2035 baseline	EM
Unitary_BT55_H_H35	High implementation of 4DT in high 2035 baseline	EM
L50_baseline	Baseline environment in 2050 with slow eco. growth and slow tech. development.	EM
Unitary_BEO1_LL_L50	Very low fuel price in low 2050 baseline	EM
Unitary_BEO1_L_L50	Low fuel price in low 2050 baseline	EM
Unitary_BEO1_H_L50	High fuel price in low 2050 baseline	EM

2050

Scenario	Description	Run by
Unitary_BEO1_HH_L50	Very high fuel price in low 2050 baseline	EM
Unitary_BTS5_L_L50	Low implementation of 4DT in low 2050 baseline	EM
Unitary_BTS5_H_L50	High implementation of 4DT in low 2050 baseline	EM
H50_baseline	Baseline environment in 2050 with fast eco. growth and fast tech. development.	EM
Unitary_BEO1_LL_H50	Very low fuel price in high 2050 baseline	EM
Unitary_BEO1_L_H50	Low fuel price in high 2050 baseline	EM
Unitary_BEO1_H_H50	High fuel price in high 2050 baseline	EM
Unitary_BEO1_HH_H50	Very high fuel price in high 2050 baseline	EM
Unitary_BTS5_L_H50	Low implementation of 4DT in high 2050 baseline	EM
Unitary_BTS5_H_H50	High implementation of 4DT in high 2050 baseline	EM

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 compile 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 the Annex.

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 for the factors correspond to time-based operations (Step1), 'Medium' to trajectory-based operations (Step2) and 'High' to performance-based operations (Step3). Based on the background scenarios we define in D3.1, we thus have the correspondence shown in Table 18.

Table 18. 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 different blocks use this information differently by changing parameters in their respective models based on their value, as described in [10], or more explicitly by implementing a different mechanism based on the underlying operational improvements.

4.3 Strategic layer

4.3.1 Economic model

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 19 summarises the effects of KPIs on the model.

Table 19. Impact of SESAR KPIs on inner parameters of the economic model

KPI	KPI description	Model parameter affected	Parameter Description
APT_CAP	Airport capacity increase	$C \rightarrow (1 + x)C$	Empirical capacity coefficient
ASP_CAP_ER, ASP_CAP_TMA	Airspace capacity increase, en-route and TMA	$A \rightarrow e^{-x}A$	ANSP pre-factor in delay generation function
PRE	Predictability increase in airport departures	<i>NI</i>	<i>NI</i>
ENV	Geometrical inefficiency decrease of trajectory	$b \rightarrow (1 + x)b$	Linear coefficient for cost of fuel
EFF_TECH	ATCO inefficiency decrease in terms of cost per ATCO	$\omega \rightarrow (1 + x)\omega$	Cost per unit of empirical capacity
EFF_ATCO, EFF	ATCO efficiency increase in terms of number of ATCO per kilometre flown	$\omega \rightarrow \omega/(1 + x)$	Cost per unit of empirical capacity
PUN	Punctuality increase in terms of number of flights below 3 minutes of delay	$\delta t_0 \rightarrow (1 - x)\delta t_0$	Airport delay generated at zero traffic

The economic model for now considers almost only the BTS factors when building the background factors. This is something which will be improved in the next version (see Section 6.2). The only other background factor used in economic model is 'BED1', which is the economic growth of the countries. For this, we used the projection from STATFOR [8] in terms of economic growth, with the different scenarios they define (fragmenting world, regulated growth and global growth, respectively applied in the Low, Medium and High scenarios). For now, we applied an average growth for the ECAC to all countries, as summarised in Table 20.

Table 20. 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%

Note that this will be improved in the next version of the model, since we have access to the growth of individual countries. In Table 21 we show how this factor and BEO1 impact the model. In terms of foreground factors, we considered BEO1, BTS5, and BTS9.

Table 21. 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

4.4 Pre-tactical layer

Two factors are modelled 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.

4.4.1 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.

4.4.2 CRCO charges (BEO2)

With the full Vista model execution, the value of the charges due to the en-route airspace service will be determined by the ANSPs based on the outcome of the economic model from the strategic layer. This factor will then determine how the charges are applied. For calibration of the model 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.

4.4.3 Factor and parameters summary

Table 22 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 22. 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.	<p>Higher fuel values will lead to the preference of more direct and shorter routes even if the total CRCO cost is increased.</p> <p>Low price of fuel will have the opposite impact, encouraging airlines to select longer routes that use cheaper airspaces when possible.</p> <p>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.</p>	<p>LL: 0.1 EUR/kg</p> <p>L: 0.2 EUR/kg</p> <p>M:0.5 EUR/kg</p> <p>H: 1 EUR/kg</p> <p>HH: 2 EUR/kg</p>

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.	<p>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.</p> <p>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.</p>	<p>Current: CRCO charges are computed as nowadays (See Section 2.2.1.3).</p> <p>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 its 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.</p> <p>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.</p> <p>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.</p> <p>Double: All rates of EUROCONTROL's CRCO's member states rates are double. This increase the relevance of en-route airspace costs.</p> <p>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.</p>

4.5 Tactical layer

For initial calibration of the tactical layer a set of three factors have been modelled. One from ‘Demand and other economics factors’ category and two from ‘Operational and technological business factors’ category. The factors affect several aspects of the tactical model: airports, airspace and passengers management which in turn demonstrates the flexibility of the tactical layer. In other to show the capabilities of the model and the effect of those factors on the overall system, two sets of parameters for each scenario have been chosen. Note that the precise values of the parameters are extracted from the expected effects described in Section 3 of D2.1. Nevertheless, the aim of this factor analysis is calibration and evaluate the modelling capabilities of the tactical layer.

As the definition of each factor is rather high conceptual, it can be interpreted in many different ways, only the impact of changing the model parameters are assessed. The tactical layer is a what-if assessment tool in the sense it is subdued to the parameters. More interesting is the intricate effects 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 to the baseline results for calibration and impact assessment. Once the baseline and factors are calibrated, all combination of two and three factors are considered. The interactions are far more complex when considering multiple factors simultaneously, calibration is no longer possible but the results allow us to explore synergies and opposite effects of the different factors.

4.5.1 Fuel prices (BEO1)

Volatility of fuel prices drives the tactical behaviour of the airspace users. The 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, on a low fuel price scenario more delay is likely to be recovered. The expected effect on the tactical model is therefore focus 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 index 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 costs. This factor is modelled thought two parameters, the maximum % of delay to be recovered as a function of the nominal route length and the amount of residual delay.

Parameter set A (table below) focus on a scenario of higher fuel prices, while set B focus on lower fuel prices. Note that the aim of the strategic layer is not to determine the effect of fuel prices on the parameters, but rather the effect of the parameters changes on the overall network metrics, both for passengers and flights.

4.5.2 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 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 us to apply mitigation strategies at local and network levels.

In the tactical model this operational and technological changes affect both the airport and airspace capacity as well as increased predictability of the gate-to-gate phase. Each of these elements is modelled differently in the tactical layer and therefore it 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 applying 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 conditional probability the chance of a given delay being ATFM is determined and in that case the amount of delay is decreased by the expected impact. Final en-route predictability is modelled by adding random noise to the nominal duration of a flight, in this context an increase on predictability is simply a reduction on the standard deviation of the random noise applied.

4.5.3 Passenger management tools (BTO4)

Improved passengers reaccommodation tools will affect how passengers are reaccommodated, in general according to a better understanding of cost of delay. As more factors are considered actual waiting times and missing connections could be either increased or decreased as not simply the delay but the costs will drive the processes.

Currently stranded passengers in the tactical layer are reaccommodated based on their ticket category and the airline archetype. There are three airline operator archetypes considered in Mercury, the classification is given in the input data set as described in Section 3.3:

- Full-service carriers, hub-and-spoke structure selling multiple flight tickets to passengers, prioritising connections;
- Regional carriers, less hub focused structure, focusing on point-to-point performance;
- Low cost and charter carriers, not offering connecting flights to passengers.

In general, full-service carriers will only reaccommodate premium passengers into other alliances 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 thumbs are modelled. However, in a future scenario, delay of some key flights could be increased to wait for premium passengers.

4.5.4 Factor and parameters summary

Table 23 shows the summary of the factors modelled in the tactical layer with information on how the factors affect the model and their expected impact. In this case, all the factors are affecting the Flight Plan Generator block.

Table 23. Summary of factors modelled in tactical layer

Factor	Summary description	Expected effect	Model affected	Implementation of different values		
				Baseline	Sample parameter-set A	Sample parameter-set B
BEO1 – Fuel prices	Fuel prices have an impact on the procedures that affect fuel consumption such as tactical delay management	Impact on operational costs and efficiency of some mechanism (e.g. delay management and recovery) Fuel consumption model update, e.g. increased delay recovered by pilots.	Flight variability	Recover (expected) arrival delay up to 5% of the route nominal duration, leaving a residual of 5 minutes.	Recover (expected) arrival delay up to 4% of the route nominal duration, leaving a residual of 7 minutes	Recover (expected) arrival delay up to 6% of the route nominal duration, leaving a residual of 3 minutes
BTS5 – 4D trajectories	4D trajectory management impact capacity and predictability	Impact on airport Capacity increased by a +0.32% En-route airspace Capacity increased by a +2% and TMA capacity by +3.62% Predictability, in terms of flight duration variability decrease, up to 4%	Airport capacity	Declared capacity (i.e. maximum number of movements per hour) is computed from traffic. Runway occupancy time is the inverse of declared capacity.	Same	Multiply baseline runway occupancy times by 1/1.032 (i.e. 0.32% increase in capacity)
			Airspace capacity	ATFM delay modelled using exponential distribution for 5% of the flights.	ATFM delay modelled using exponential distribution for 4.5% of the flights.	ATFM delay modelled using exponential distribution for 3% of the flights.
			Flight variability	On-route uncertainty modelled as a normal distribution of zero mean and standard deviation of 2% of the route nominal duration.	Use a standard deviation of 1.88% instead (i.e. $2*(1-0.97)^2$)	Use a standard deviation of 1.84% instead (i.e. $2*(1-0.96)^2$)

Factor	Summary description	Expected effect	Model affected	Implementation of different values		
				Baseline	Sample parameter-set A	Sample parameter-set B
BTO4 – Passengers tools	Passengers’ re-accommodation tools affect passengers’ re-accommodation procedures and tactical delay recovery strategies as it impact cost of delay.	Reduction of operational costs and increased performance on passenger delay management.	Passenger re-accommodation	Re-accommodation priority:	Passengers are re-accommodated in the next available flight, regardless of the AO:	Passengers are re-accommodated in the next available flight, regardless of the AO:
				- NON FSC tickets: only same airline - FSC standard tickets: less than 5 hours, only same alliance, else all - FSC premium tickets: all possible flight		
			Passenger waiting rules	If flight is not ATFM regulated, wait up to 10 minutes until the flight has 80% of expected pax	If flight is not ATFM regulated, wait up to 30 minutes until the flight has 80% of expected pax for either one premium connecting pax, two onward long-haul passengers	If flight is not ATFM regulated, wait up to 60 minutes until the flight has 80% of expected pax for either one premium connecting pax, two onward long-haul passengers

5 Initial results

5.1 Strategic layer

The scenarios tested for the strategic layer aims to evaluate the impact of increased demand, of the price of fuel and the benefits from the foreseen SESAR improvements. Some of these scenarios are realised in 2035 and 2050.

Regarding metrics, we focussed in this deliverable on average measures, however it should be kept in mind that this can hide very different realities for the agents. We used several metrics per agent, as listed below.

- Airport:
 - capacity
 - (departure) delay (per flight)
 - total costs
 - profit
 - traffic (for 24 hours)
- Airline:
 - (departure) delay per OD pair
 - ATFM delay per OD pair
 - ATC cost
 - airport cost (landing/terminal fees)
 - crew cost (per pax)
 - maintenance cost (per pax)
 - fuel cost (per pax)
 - capital cost (per pax)
 - ticket price
 - fuel consumption (per OD pair)
 - profit (per flight)
 - number of passengers (per OD pair)

- ANSP:
 - capacity
 - delay
 - costs
 - revenues
 - profit
 - unit rate
- Pax (all per itinerary):
 - airport delay
 - ATFM delay
 - ticket price
 - frequency
 - volume

Due to the number of metrics, displaying and analysing them is a challenge. In the following, we try to highlight the main results.

5.1.1 Evolution of baselines in 2035 and 2050

We begin by showing the evolution of a set of metrics in the baseline scenarios for 2035 and 2050.

In Figure 44, we first show the time evolution of a few key metrics from airlines, airports, ANSPs, and passengers for the two Low and High baseline scenarios.

Combined metrics

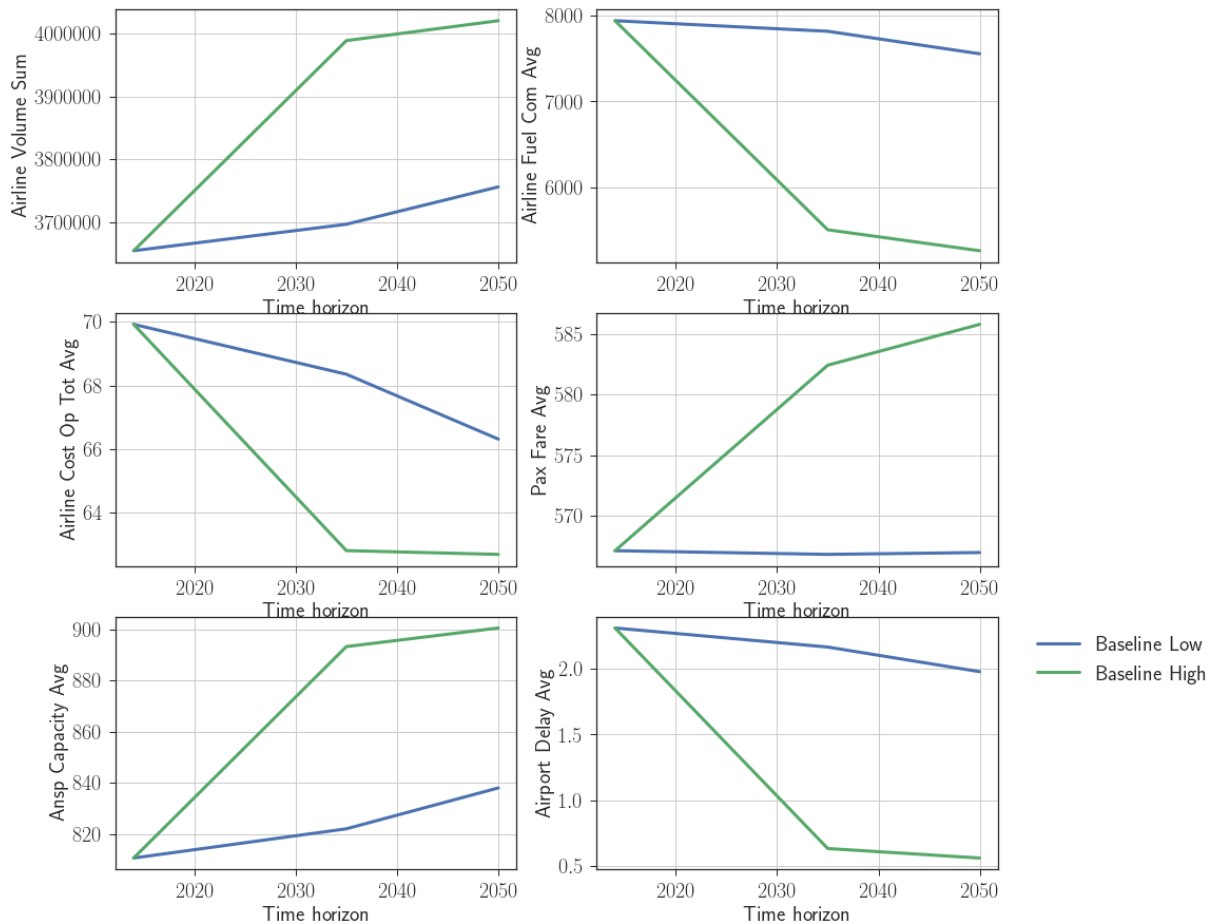


Figure 44. Six key metrics for different stakeholders in the baseline scenarios

Baseline scenarios: Current, Low 35, High 35, Low 50, High 50. From top to bottom, and left to right: total traffic in passengers, average fuel consumption in kg, average operational cost for airlines (in euros), average ticket price (in euros), average ANSP capacity, and average airport departure delay (in minutes).

First, we note that even in the High baseline scenario, the increase of traffic overall in traffic seems quite moderate (9% of growth maximum in 2050 w.r.t the current situation). This is in sharp contrast with the STATFOR prediction (around 175% increase for the most optimistic scenario), and might indicate a calibration issue. In particular, the demand might be under-sensitive to changes in income. We also did not take into account population growth, but only the GDP per capita to compute the increase. The fuel consumption per flight also falls, and as we will see in the following, this is not counterbalanced by the increase in traffic, so that the emissions decrease overall, especially in the High scenario. The average operational cost for airlines falls significantly too, which is partly translated into lower prices. Both airport delay and ANSP capacity also decrease.

This representation is useful to see the evolution of a few metrics but does not allow us to have a complete picture. For this reason, in Figure 45 we show another representation where all metrics are displayed at once, for a given time baseline (Low here).




Figure 45. Airport metrics in the baseline scenarios

From left to right: average capacity, average departure delay; average total costs (expenses) per pax, average profit per pax, average revenues per pax (all in euros); average traffic (number of passengers in a day).

Regarding airports, the first remark is that they are really passive in the model because we deactivated their capacity extension. Moreover, the above figure is computed based on averages over all the airports, including non-European airports, which means that some values are unexpectedly low (for instance for traffic, European passengers in non-European airports were included). However, some trends are clear, and due to various SESAR improvements, their capacity has significantly improved. Since the traffic increase is not highly consequential, even in 2050 the airports’ average delay drops (conversely), albeit by a small amount. Airport costs increase substantially, since we did not change the cost of capacity, but increased their actual capacity. This is due to the way we interpreted the SESAR KPI ‘Capacity’, which says nothing about the cost of the increase, or, equivalently, the efficiency of delivering the capacity. This is contrast with ANSPs (see below for more discussion), which have a separate target for their capacity and for their efficiency. Finally, as expected, airport revenues increase, but not enough to counterbalance their loss.




Figure 46. Airline metrics in baseline scenarios

From left to right: the average departure delay per flight, the average ATFM delay per flight, the average ATC cost per pax, the average capital cost per pax, the average aeronautical fees at airport per pax, the average crew cost per pax, the average strategic cost of delay per pax, the average fixed cost (not used), the average cost of fuel per pax, the average maintenance cost per pax, the average fare on an OD pair, the average fuel consumption per flight, the average profit per flight, and the average number of pax per OD pair.




Figure 47. ANSP metrics in baseline scenarios

From left to right: average flight-kilometre flown times the unit rates (not used here), the average capacity, the average ATFM delay, the average total costs (expenses), the average total revenues, the average traffic (number of flights), the average unit rate (in cents of euros).

Moving to airlines (Figure 46), many of their metrics are not impacted by the different time horizons. Most notably, the cost of ATFM and airport delay decrease (because of gains in efficiency/capacity for both), as well as their fuel consumption (due to better trajectory efficiency). This, combined with a modest increase of traffic, allows them to be more profitable, and also decreases the price of tickets, although only modestly.

Metrics for ANSPs are displayed in Figure 47. First, we see that ANSPs seem to be a little ‘too good’ regarding delay, which does not vary from their target. Their gain in efficiency and capacity, however, allows them to deliver the same level of service with largely reduced costs. As a consequence, forcing them to be profit-neutral impacts directly the unit rate, which, on average, decreases by almost a factor of 4. This behaviour is probably unrealistic, but arises directly from the SESAR targets set for 2035 and beyond, where efficiency is expected to increase substantially. There is also a problem of interpretation of the targets, since it is not clear, as for the airports, what an increase of capacity means in practice. For instance, does it mean the ANSP is able to deliver a higher capacity, keeping its costs fixed? Moreover, it is clear that for ANSPs the maximum instantaneous capacity they can provide is capped. This means that the effective capacity (as we have defined it in Section 2.1.1) is somewhat unclear and needs refining. This is listed in Section 6.2 as potential future improvement of the model.




Figure 48. Pax metrics in baselines scenarios

From left to right: Average delay per itinerary, average ATFM delay per itinerary, average fare per itinerary, average frequency per itinerary, average number of passengers per itinerary.

Finally, we show in Figure 48 the (currently) more limited pax metrics (see also Section 6.2 regarding extension of the model to the door-to-door context). Most of these are readily understandable given

the discussion above. An interesting point, however, is that the frequency of itineraries is quite constant across different time horizons. Note also that the fare does not coincide with the above ticket price for airlines because here it represents the price of the complete itinerary, whereas the figure for airlines showed the price for single legs.

5.1.2 Impact of fuel price

The price of fuel is one of the most important business driver for airlines, since, as we will see in the following, its cost represents an important share of the total operational cost of the airline. It is important to realise that it also has different impacts on different airlines. Because some airlines have a higher profit margin, they can more easily keep their operation running even if the price of fuel increases substantially. They can even capture some traffic from other airlines with smaller profit margins, and thus benefit overall from the price increase. This can be demonstrated by the model, as has been shown in an earlier toy version.

First, we show in Figure 49 a representation of the airlines’ metrics in the Current scenario, the L35 baseline and both L35 with low and high fuel price. Some of these metrics vary in a very intuitive way. For instance, the cost of fuel is much higher when the price of fuel is high, and smaller when the price decreases, with respect to the baseline. Airlines make more money (higher profit) when the price is low, and the traffic increases, driven also by lower prices.




Figure 49. Airline metrics by scenario

For explanation of the metrics, see caption of Figure 46

However, other metrics vary counterintuitively. For instance, fuel consumption is smaller both in the low price and the high price scenario compared to the baseline. Similarly, ATC costs and (the cost of) delay both decrease in these two scenarios compared to the baseline. The reason behind these behaviours are not yet clear. One explanation is erroneous code. Another explanation is related to

the unexpected behaviour of the ANSPs, which have vastly improved efficiency and sometimes drive the airlines to use more efficient trajectories (hence the smaller cost of delay and fuel consumption). A third possibility is that this is an emergent behaviour due to the heterogeneity in the system. For instance, it is possible that decreasing the fuel price benefits more a certain type of airline, typically with a smaller profit margin (e.g. a LCC). These airlines could have a better fuel consumption profile, because they operate mainly short-haul routes, which in turns drives the average fuel consumption down. This is consistent with the fact that the average delay at airports also decreases: the average delay is expected to decrease as traffic shifts from highly congested to secondary airports, mostly operated by LCCs.

Because the system is highly non-linear, increasing the fuel price does not have the same effect, necessarily. For instance, because airlines all have small margin profits when the price is ‘too high’, and thus have more homogeneous behaviours, they may all decrease their traffic in the same way to decrease their fuel consumption. This should become clearer in the next iteration of the model.

The issue of the cost structure of airlines is central to the model, and can be considered as one of its main features: airlines, even though they are built with the **exact same behavioural rules**, can *in fine* display different characteristics because of their heterogeneity in terms of parameters. The cost structure of the airline in the different scenarios is displayed in Figure 50, to show how they evolve through time and scenarios. **We believe that this cost structure, more than any other factors, will play a crucial role in the future air transportation landscape.**

Operational cost structure of the airlines




Figure 50. Operational cost structure of the airlines by scenario (percentages)

Costs considered are: ATC cost, airport charges, cost of fuel, cost of crew, and maintenance cost. Cost of delay omitted for easier comparison with empirical data.

It is first interesting to see that the repartition of the costs are quite similar in the current scenario and in the 2035, Low baseline. Only the cost of delay and ATC costs change, and only moderately. As expected, the share of the cost of fuel changes quite substantially between the low price and high price scenario, to peak at around half of the operational cost in the latter case. (Another point, the vanishing ATC costs, which are unrealistic as pointed out above, should be corrected in the next version.)

In Figure 51, we show the same kind of graph for the High scenario. Most of the conclusions are qualitatively the same, but it is interesting to see that the gain in the high scenarios are able to offset partly the high price. For instance, the traffic volume in the high price scenario is still higher than in the current scenario, whereas in the low scenario the traffic decreased below the initial baseline. The fuel consumption has roughly the same qualitative behaviour, even though this time the effect of the price is quite small compared to the effect of the baseline itself. The fares are always higher than the current scenarios, which is explained by the strong increase in demand in the High scenario.




Figure 51. Further airline metrics by scenario

For explanation of the metrics, see caption of Figure 46

The results for 2050 are very similar to those obtained for 2035, with different magnitudes. In particular, the ANSPs are so efficient that their unit rates completely vanish.

5.1.3 Impact of 4D trajectories

The second foreground factor studied in this deliverable for the strategic layer is coded 'BTS5' and corresponds to various improvements enabling full 4D trajectory management. We tested two levels of implementation of this factor, as explained in Section 4.3, and simulated them in the Low and High baselines, for 2035 and 2050.

Figure 52 shows the airline metrics for the Low scenarios, with low and high 4DT implementation for 2035. The first remark is that 4D trajectories seem to have a significant impact on the system, at least in its high implementation version. In terms of cost, implementing 4D trajectories decreases the ATC costs, as expected, and the cost of delay, because of more efficient trajectories. The fuel consumption decreases too, and, correspondingly, the airlines are able to supply more seats, which drives the traffic up.




Figure 52: Airline metrics under different scenarios, testing 4DT implementation levels

For explanation of the metrics, see caption of Figure 46

If the situation is similar for 2050, with further gains, the High scenarios tell a slightly different story, as seen in Figure 53. Indeed, this time the gains in the ATC cost are not so important compared to the baseline. The fuel consumption, profits, traffic, etc., are then intermediate between the current and baseline, for both levels of implementation. This is again counterintuitive as above, and will need more investigation to be fully understood. It is possible once again that in these scenarios, different airlines gain more from the implementation when it is low or high, thus changing their supply differently and driving the overall average in different directions.




Figure 53. Testing 4DT implementation levels for different scenarios with the H35 baseline

For explanation of the metrics, see caption of Figure 46

In order to see this effect more clearly, we show in Figure 54 the results of some key metrics for the different scenarios. It is striking that airlines have the lowest operational cost in the baseline High scenario, with nevertheless some low airport delay and high ANSP capacity. The passengers also have the highest fare in average in this scenario, whereas it is fairly constant in the Low baseline scenario.

Combined metrics

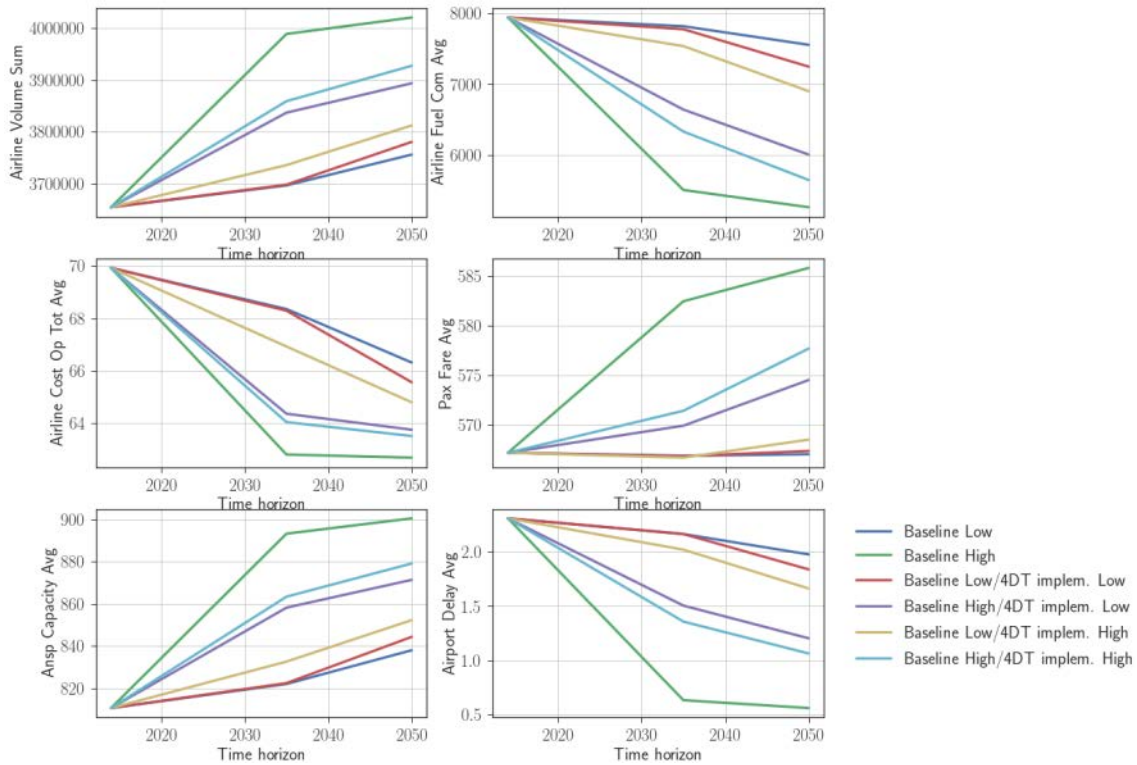


Figure 54. Time evolution of key metrics when testing 4DT implementation levels

For explanation of the metrics, see caption of Figure 46

5.1.4 High-level view for trade-off analysis

In this section, we show some high-level views for the different scenarios that we explored above, which can be useful to perform a trade-off analysis. To do this, we selected five KPIs, one per stakeholder considered in Vista:

- Average ANSP capacity;
- Average airport delay per flight;
- Average ticket fare (for pax);
- Total CO₂ emissions (for environment);
- Average total operational cost for airlines.

Baselines

Comparing the baselines with this high-level view is useful to have an intuition on what is the general direction of the system. On the left of Figure 55, we show the evolution in the Low baseline scenario. In this case, the system is pretty stable, with increased ANSP capacities, decreased airport delay, decreased airline costs, and decreased total CO₂ emissions. The picture for the High baseline scenario on the right is similar, with most gains. It is interesting to see that most of the gains are made in the first phase (until 2035), and that the KPIs are fairly stable after that. It is also interesting to see how the average fare increases a little, undoubtedly because of the step increase in income in this scenario.




Figure 55. 'Radar' plots for the 5 KPIs in the baseline scenarios

Left: Low baseline. Right: High baseline. Values are normalised by the highest ones for each KPI.

Fuel price

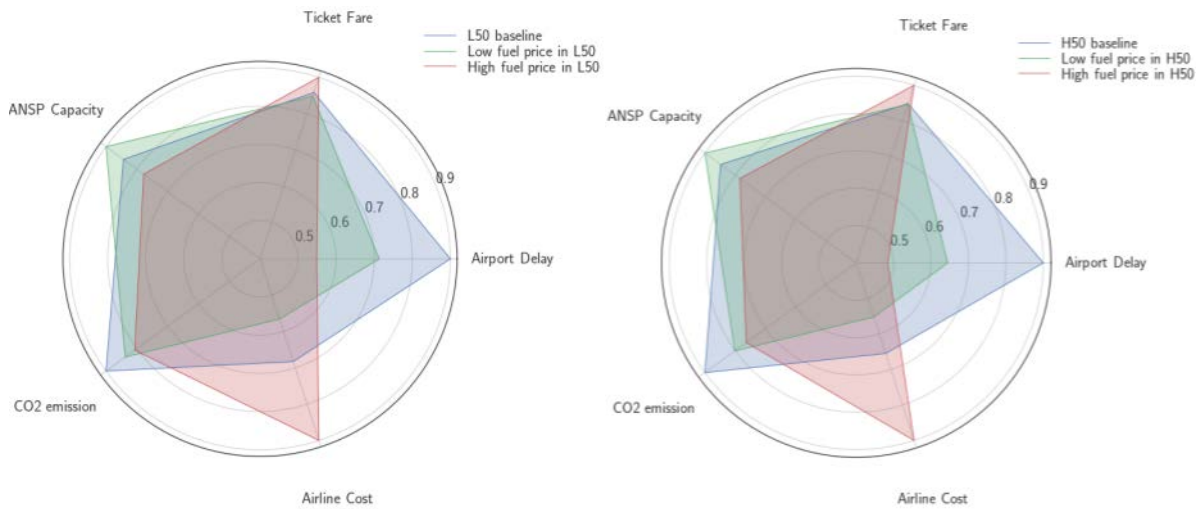


Figure 56. ‘Radar’ plots for the 5 KPIs when changing price of fuel in 2050

In Figure 56, we show what happens with the change of price with the Low and High baselines in 2050 (plots for 2035 are very similar, with more moderate changes). As we have seen before, changing the price of fuel has a non-trivial effect on the system. The ANSP capacity and airport delay are monotonic with the price of fuel, which tends to decrease the ANSP capacity – the agents adapt to lower traffic conditions – and decreases also the airport delay – because the airports do not adapt their capacity. The ticket fare also is slightly smaller for the lower fuel price, and higher for the higher fuel price. Note that the degree to which the increase in fuel price is passed on to the ticket price seems non-linear. In other words, airlines internalise more or less their loss/gain in this regard. In this regard, the model is able to produce quantitative measurements to weigh in the long-standing debate regarding whether or not airlines internalise some of their costs. The fuel consumption overall is smaller for both low and prices, which could be due to the heterogeneity of airlines, as discussed above. The total cost for airlines increases with the price of fuel.

It is interesting to see that the conclusions are roughly the same for the Low and High scenarios, which shows how important the price of fuel is with respect to the overall operational situation.

4D trajectories implementation

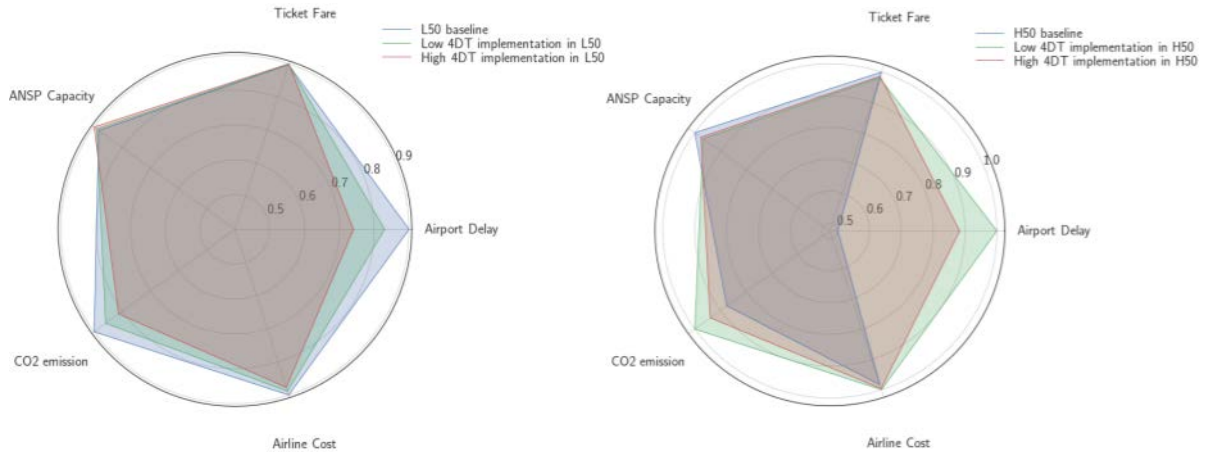


Figure 57. ‘Radar’ plots for the 5 KPIs in when changing the 4DT implementation level in 2050

Finally, we show in Figure 57 how the 4DT implementation changes the KPIs. Overall, the impact seems positive for all stakeholders, according to these KPIs. The ANSP capacity and ticket fare are quite stable, but the airport delay, and the emissions drops by a sizeable amount (10% max for CO₂, almost 30% for delay). Note that this is not contradictory with the conclusions drawn above in Section 5.1.3, since the latter were based on the average fuel consumption per flight, whereas here we compute the total emissions. Finally, the overall costs for the airline seem quite stable, with a slight decrease when implementing 4DT.

5.2 Pre-tactical layer

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 different scenarios described in Section 4 have been executed. 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.

5.2.1 Impact of fuel cost on flight plan selection




Figure 58. Impact of fuel cost on FP selection characteristics

Figure 58 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 59. Impact of fuel cost on FP selection characteristics; flights with one route option




Figure 60. 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 59). 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 60).

5.2.2 Impact of en-route airspace charges on flight plan selection




Figure 61. Impact of en-route charges on flight plan selection characteristics

Figure 61 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 than 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 flight 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 flights 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 5.2.4.

5.2.3 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 62. En-route airspace usage revenues per NAS for different fuel scenarios

Note AZ: Portugal Azores airspace, EX: Shanwick airspace

Figure 62 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.

5.2.4 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 63. Impact of fuel cost on FP selection characteristics for GCTS to EGKK flights




Figure 64. 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 69.a to see all possible route between those two regions).

Figure 63 and Figure 64 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 5.2.2, 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 65. Impact of en-route charges on FP selection for GCTS to EGKK flights


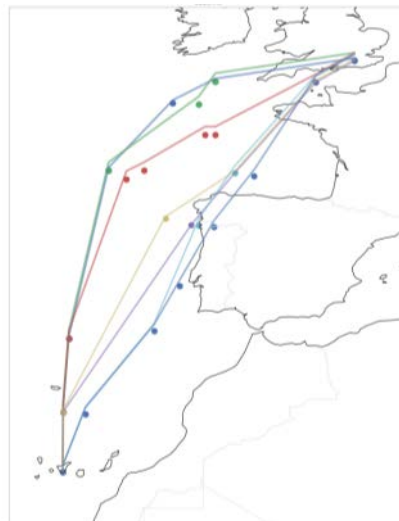



Figure 66. Impact of en-route charges on FP selection for all EG–GC flights


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 65 and Figure 66).




a) Baseline scenario




b) Very low fuel



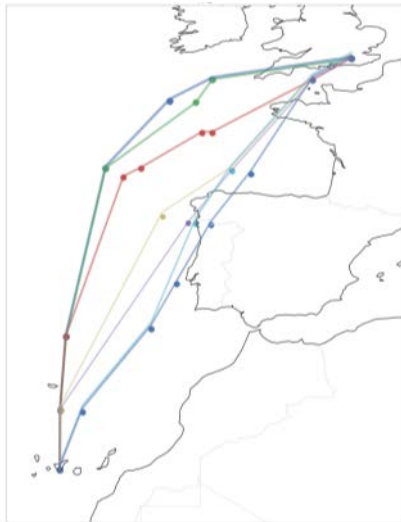
c) Low fuel



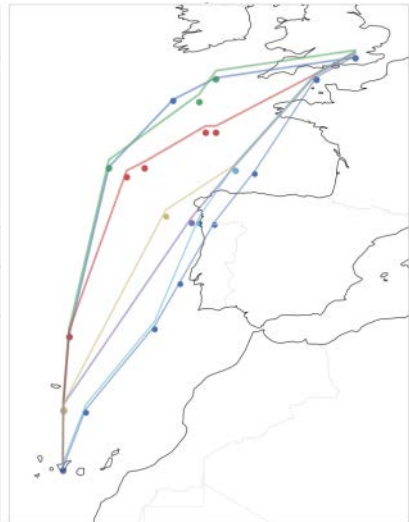
d) High fuel



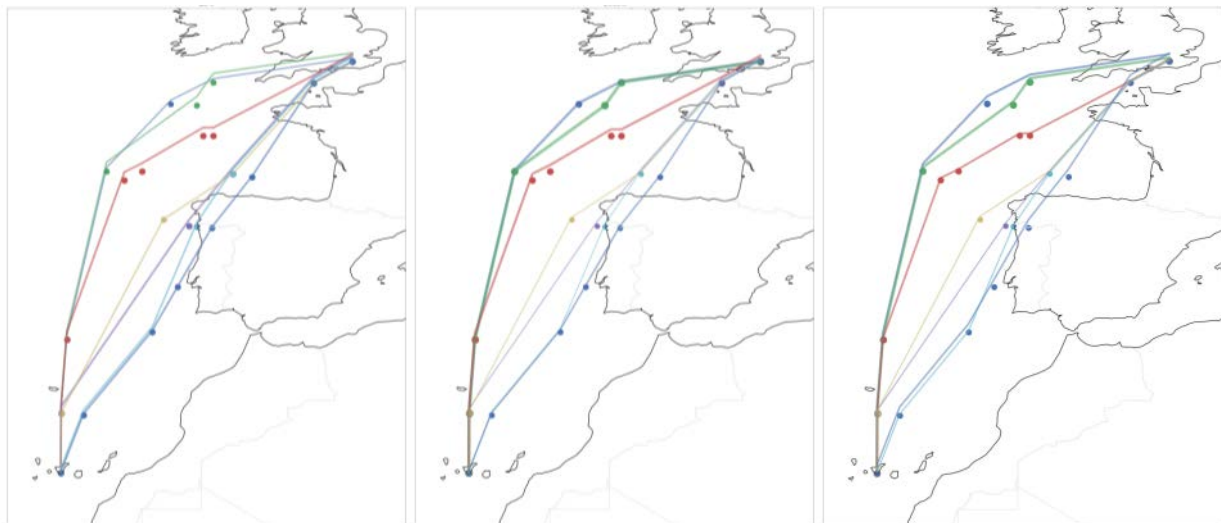
e) Very high fuel



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 67. Routes selected by flights between EGKK to GCTS for different scenarios

Figure 67 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 not very big, but for example, comparing Figure 67.g (homogenous airspace charges) with Figure 67.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 68. 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.


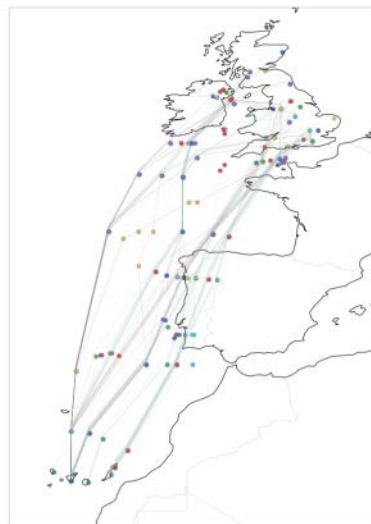
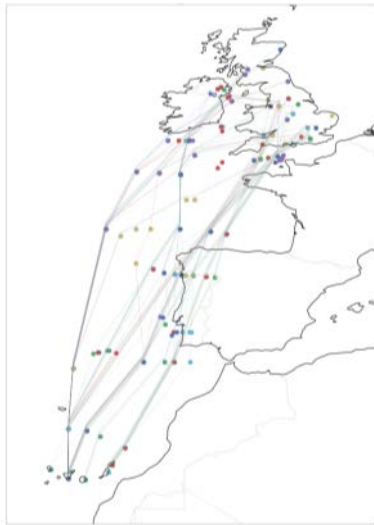



Figure 68. Demand and revenue variations, EGKK–GCTS, when changes in fuel price produced



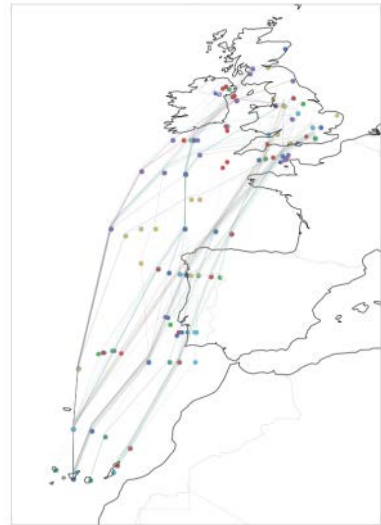
a) Baseline



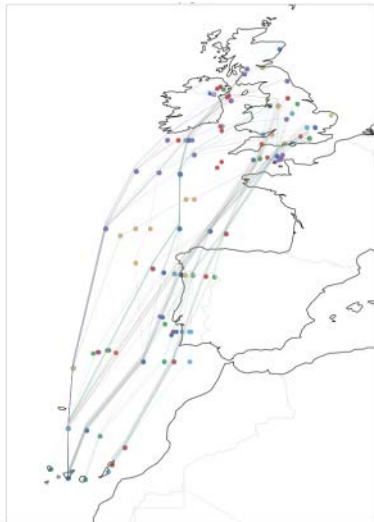
b) Very low fuel




c) Low fuel



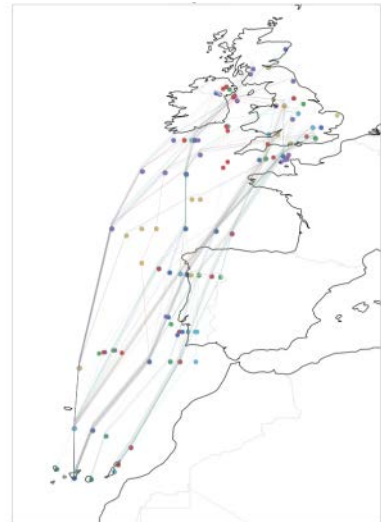
d) High fuel



e) Very high fuel



f) FABs



g) Homogeneous

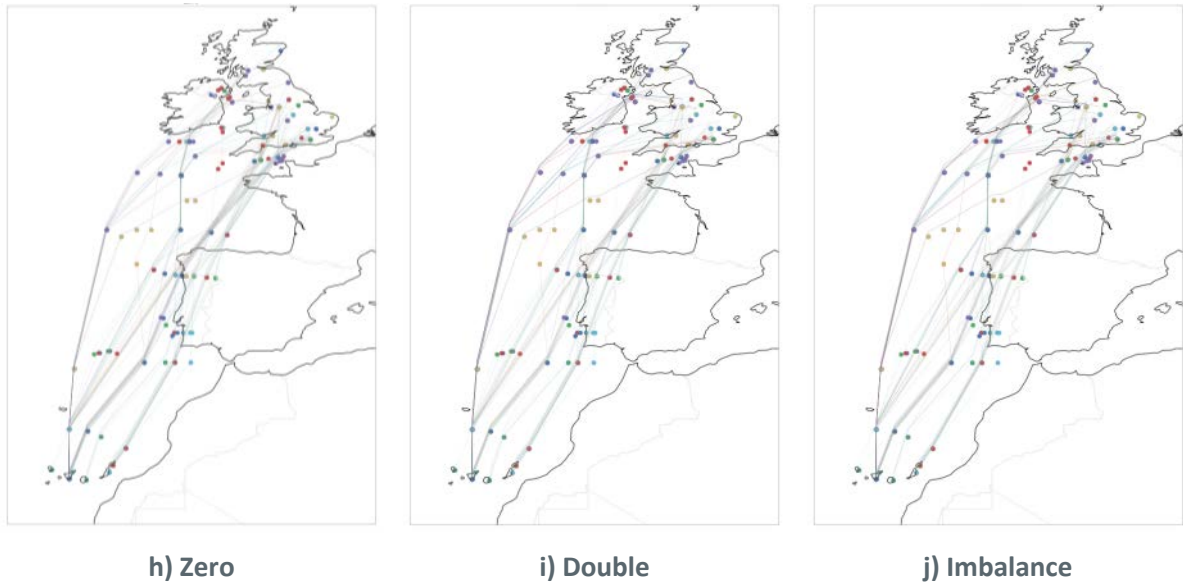


Figure 69. Routes selected by flights EG to GC for different scenarios




Figure 70. Demand and revenue variation, EG to GC, when changes in fuel price produced

When all the flights between EG and GC are considered (Figure 69), the variations on demand and revenues due to changes in fuel price are as expected (Figure 70) 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%.

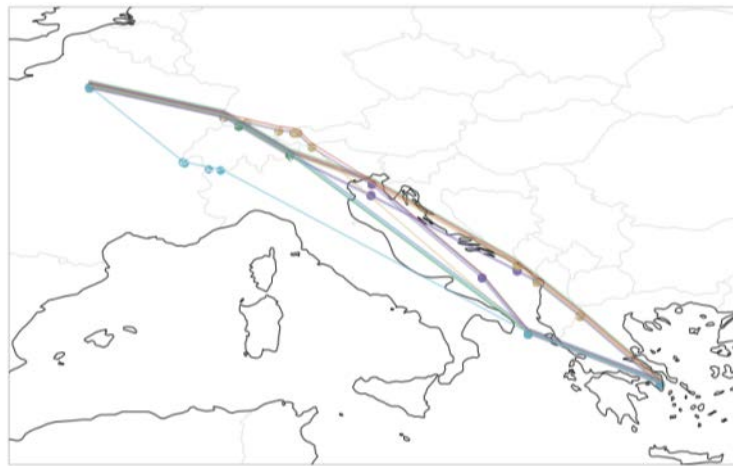


Figure 71. Possible routes between LPFG to LGAV as selected in baseline scenario

Figure 71 presents routes selected for flights between LPFG 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 72 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 73 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.




Figure 72. Demand and revenue variation, LPGA to LGAV, with changes in fuel price




Figure 73. Demand variation LPGA to LGAV with changes in airspace charges (and fuel prices)

5.2.5 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 74 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.




Figure 74. Aircraft entries in 30-min time windows by NAS and fuel scenario

5.3 Tactical layer

Table 24 summarises the reference metrics for the preliminary analysis of the factors and scenarios defined in Section 4.1. Baseline reference values are obtained when no factors are applied, i.e. current air transportation system. Metrics are focused on delay, both for passengers and flights which give a limited but sufficient for calibration.

Lastly, note that the aim of the tactical layer is to serve as a what-if tool that assesses the impact of the different factors and detect trade-offs, synergies and opposites. Therefore, reference values should be considered as a base framework to normalise results. No metric should be considered as an absolute value, but relatively.

Table 24. Tactical metrics computed and baseline reference value

Name	Formula	Baseline reference value
Average flight departure delay	$\frac{\sum_f (AIBT_f - SIBT_f)}{\sum f}$ <p>where f ranges over all not excluded and not cancelled flights</p>	4.65 min/flight
Average flight arrival delay	$\frac{\sum_f (AOBT_f - SOBT_f)}{\sum f}$ <p>where f ranges over all not excluded and not cancelled flights</p>	5.99 min/flight
Percentage of reactionary delay	$\frac{\sum_f (AOBT_f - SOBT_f)}{\sum_g (AOBT_g - SOBT_g)}$ <p>where f ranges over all not excluded and not cancelled flight delayed by either aircraft or passenger or both, and g ranges over all not excluded and not cancelled flights</p>	40%
Average Passenger departure delay	$\frac{\sum_f (AOBT_{f'_o(p)} - SOBT_{f_o(p)})}{\sum p}$ <p>where p ranges over all passengers $f_o(p)$ and $f'_o(p)$ are the originally booked and actual first leg of the passenger's itinerary</p>	4.27 min/pax
Average Passenger arrival delay	$\frac{\sum_p (AIBT_{f'_N(p)} - SIBT_{f_N(p)})}{\sum p}$ <p>where p ranges over all passengers $f_N(p)$ and $f'_N(p)$ are the originally booked and actual last leg of the passenger's itinerary</p>	5.88 min/pax

5.3.1 Isolated factor results

As expected, isolated factors have a limited impact on the global system that depends heavily on the exact parameters (see Figure 75). For instance, 4D trajectories (BTS5) improves both passenger and flight delay metrics, however it also increases reactionary delay ratio in both parameters sets. This shows that increased reactionary delay could be a viable strategy in some scenarios, at least from the delay perspective.




Figure 75. Initial results of the tactical layer for isolated factors

On the contrary better passenger tools reaccommodation tools (BTO) as implemented in Section 0 do not show a significant improvement of the metrics considered. Just measuring delay is not enough to capture the effect of this factor, it will be more illustrative to explore more passenger centred metric such that the number of missing connections.

Finally, results show that fuel prices (BEO) is the more relevant factor of those considered. The high fuel cost scenario, implies less delay recovered that results in higher arrival delays, higher reactionary rate and therefore also larger departure delays. The resilience of the system is reduced greatly when fuel prices increase, this reduces the availability of the system to recover delay. On the contrary, when fuel prices decrease, more delay can be recovered and the resilience of the system increases leading to better delay metrics, both for passengers and flights. Note that fuel is not currently explicitly model and hence fuel cost has an impact on the rules used to recover delay but these are not cost driven.

5.3.2 Cross-factor results

More interesting is the combined results of two or more factors applied simultaneously (see Figure 76). The best synergy appears when 4D routes (BTS5) are applied on a low fuel price scenario (BEO1 and parameter set A) it is a win-win situation that increases all metrics considered. In further analysis this synergies and win-win situations will be explored in more details. It is possible that this situation

is true only for some metrics or areas of performance and that there is an actual trade-off between different areas.




Figure 76. Initial results of the tactical layer for paired factors

Although not as good as previously mentioned, there is also a synergy between better passenger reaccommodation (BTO4) tools and 4D routes (BTS5). Delay is overall decreased, except the reactionary delay ratio which is actually increased. Nevertheless, reactionary delay ratio is just an indicator of delay sources and not directly performance. It is also interesting that some scenarios compensate, e.g., in scenario B when BEO1 and BTO4 are active and scenario A when BTS5 and BTO4 are active simultaneously. In these scenarios the negative impact of some factors, e.g., higher fuel prices (BEO), are compensated with the positive impact of other factors, e.g. 4D trajectories (BTS5).

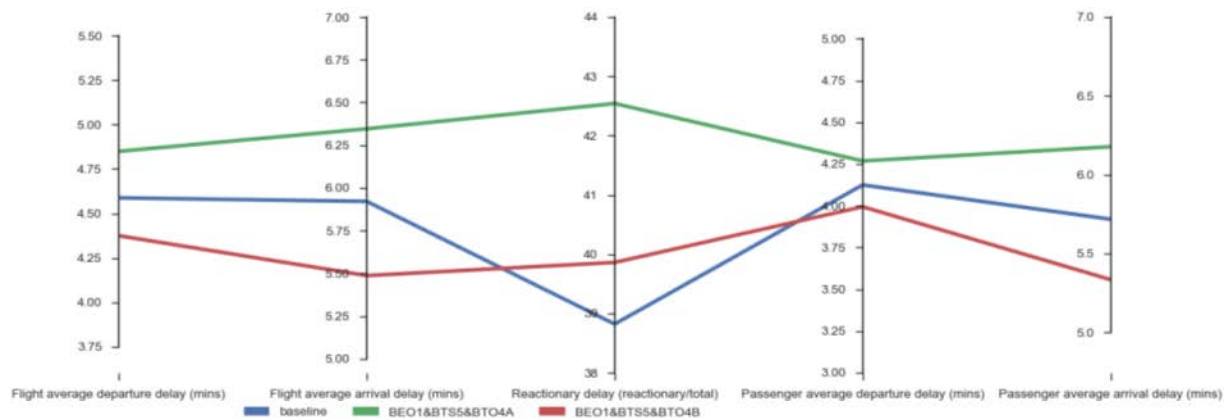


Figure 77. Initial results of the tactical layer when the three factors simultaneously active

Finally, when all three factors are applied it is clear that fuel prices (BEO1) still dominates (see Figure 77). The effect of fuel prices may be partially compensated by the other factors, but still higher prices imply worsening of delay metrics and *vice versa*. It is also interesting to realise that for several factors, despite increasing reactionary delay it does not imply worst departure and arrival metrics overall, for both passengers and flights.

6 Next steps and look ahead

6.1 Model layer integration and consultation in progress

This deliverable has demonstrated the capabilities of each of the different layers of the Vista model: strategic, pre-tactical and tactical. Vista is already able to model complex scenarios with single or multiple factors and produce results from each of the layers. The effort devoted by the team to the calibration has ensured that the individual layers are working as expected. The Vista team will carry out a consultation on these initial results in order to enhance some of the modelling assumptions and to identify the type of analysis that should be performed on the final results. Part of this consultation has already started with the presentation of some of the initial results in the PJ19.04 Progress Meeting in Madrid on 06MAR18. The outcome of these consultation activities will be reported in D6.3 (due in MAR18). In the coming weeks, the different layers will be integrated allowing the execution of a given scenario by the three layers. Further development of the model is also foreseen, as described in the next subsection.

6.2 Features to be potentially deployed in final model version

This is a list of features to be potentially included in the final version of the model (reported in Deliverable 5.2). This list will be prioritised internally, also using feedback from the consultation process (reported in D6.3).

- Implement the gate-to-door/door-to-gate features of the model
 - *These are being finalised in parallel and will be included in the next (final) version of the model, along with limited corresponding KPIs (it is important, as previously discussed, to limit the number of indicators analysed, to meaningfully focus on the main effects and model outputs)*
- Explicitly flag airports that can extend their capacity
 - *A preliminary attempt has been made in this direction, although it is a highly demanding task; all capacity extensions in the model are currently frozen*
- Develop a refined set of rules for future ANSP behaviour
- Model improved minimum aircraft turnaround times
- Include business and non-business pax in the *economic* model
- Coordination with the ATM4E project on the inclusion of NO_x
 - *A productive webex with the SJU and DLR was held in early MAR18, on the basis of which the Vista team is now developing the NO_x model components*
- Integrate more background factors into the models
- Use more parameters on the flight plan selection, particularly temporal and buffer-related
- Development of the Itinerary Generator and its calibration with historical data
- Develop an ATFM functionality dependent on demand
- Selection of flight plans at the tactical level, as a function of the tactical situation (e.g. delays)
- Further develop the application of GDP growth per country
- Further develop the post-calibration procedures for the economic model

7 References

1. 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.
2. IATA, "Air travel demand", IATA Economics Briefing no 9, 2008.
3. Adler, N., "Competition in a deregulated air transportation market", *European Journal of Operational Research* 129, 337-345, 2001.
4. DATASET2050 project, Deliverable 3.1, "Current passenger profile", 2017.
5. Cook, A., Tanner, G. "European Airline Delay Cost Reference Values, Updated and Extended Values", version 4.1, 2015.
6. Performance Review Commission, "Performance Review Report 2016", June 2017.
7. EUROCONTROL, CODA DIGEST - All-Causes Delay and Cancellations to Air Transport in Europe – Annual 2014, 2015.
8. STATFOR, "Challenges of Growth", 2013.
9. EUROCONTROL, Monthly follow up of service Units All CRCO Charging zones, 2018.
10. SESAR JU, Vista Deliverable D3.1 "Business and Regulatory Scenarios Report", 2017.
11. EUROCONTROL, Monthly adjusted unit rates, February 2018.
12. ICAO, "ICAO DATA+", <http://www4.icao.int/newdataplus>, accessed September 2017.

8 Annex: Interpolated and extrapolated targets

Table 25. 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 26. 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 27. 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%

-END OF DOCUMENT-

