D4.2 Final model results

Deliverable ID:	**D4.2**
Dissemination Level:	**PU**
Project Acronym:	**BEACON**
Grant:	**893100**
Call:	**H2020-SESAR-2019-2**
Topic:	**SESAR-ER4-08-2019**
Consortium Coordinator:	**University of Westminster**
Edition Date:	**23 May 2022**
Edition:	**01.00.00**
Template Edition:	**02.00.02**

Authoring & Approval

Authors of the document

Name/Beneficiary	Position/Title	Date
David Mocholí / NOM	Project Team	23/05/2022

Reviewers internal to the project

Name/Beneficiary	Position/Title	Date
Gérald Gurtner / UoW	Project Coordinator	13/05/2022
Tanja Bolic / UoW	Project Team	17/05/2022

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

Name/Beneficiary	Position/Title	Date
Gérald Gurtner / UoW	Project Coordinator	23/05/2022

Rejected By - Representatives of beneficiaries involved in the project

Name/Beneficiary	Position/Title	Date
N/A		

Document History

Edition	Date	Status	Author	Justification
00.00.01	13/05/2022	Draft	David Mocholí	Initial draft for internal review
00.00.02	17/05/2022	Draft	David Mocholí	Second draft for internal review
01.00.00	23/05/2022	Draft	David Mocholí	For SJU submission

Copyright Statement

BEACON

BEHAVIOURAL ECONOMICS FOR ATM CONCEPTS

This Technical Report is part of a project that has received funding from the SESAR Joint Undertaking under grant agreement No 893100 under European Union's Horizon 2020 research and innovation programme.

Abstract

This deliverable presents the small-scale agent-based simulation model used in WP4 to analyse the performance of different flight prioritisation mechanisms selected in WP3. The document contains a detailed description of the model and its most important components. Additionally, it offers the final model results and conclusions obtained with the set of simulation experiments, using KPIs defined in WP3.

Table of Contents

List of Tables

List of Figures

1 Introduction

1.1 Scope and objectives

One of the main goals of the BEACON project is to explore how behavioural economics can complement simulation models employed in ATM. Agent-based modelling (ABM) appears as a great opportunity to meet these objectives by, on the one hand, allowing the observation of the emergent behaviour (e.g., network effects) arising from agents' interactions in a bottom-up process and, on the other hand, providing a natural framework to incorporate behavioural economics insights about human and institutional behaviour into operational simulation models.

This has been done, in a first stage, through the development of a small-scale agent-based model, built on a simplified network model with a reduced number of airports, airspace users (AUs) and airspace sectors, which has been used to analyse how different non-rationality assumptions affect the performance of the different flight prioritisation previously selected. This simulator model has been developed based on an existing small-scale agent-based model produced by Nommon in the context of a previous SESAR Engage KTN Catalyst fund project called "Exploring future UDPP concepts through computational behavioural economics".

The work described in this document provides a detailed and exhaustive description of this upgraded and extended agent-based model. This model has been used to assess the performance of the new flight prioritisation concepts, defined in WP3, under the different agent rationality assumptions calibrated in D4.1, assessing their robustness to unexpected behaviours. These results, which are also shown and discussed in detail in this deliverable, have helped to understand how the mechanisms were behaving and to fine-tune the selected mechanisms and behaviours for larger scale impact testing to be performed in WP5.

1.2 Structure of the document

The document is structured as follows:

- Section 1 introduces the document explaining its aim and scope and describes the structure of the report.
- Section 2 provides an overview on the state of the art in agent-based modelling and its applicability to the evaluation of flight prioritisation mechanisms.
- Section 3 describes the model characteristics, including an overall description, the main modelling assumptions, and all the necessary information about the main model components.
- Section 4 details how the simulation scenarios have been defined enumerating the different aspects taken into account for that characterisation.
- Section 5 presents the analysis of the simulation results organised by the different KPAs identified within the BEACON performance framework.
- Section 6 summarizes the main conclusions of the study.

2 Agent-based modelling: Applicability to prioritisation mechanisms

The majority of the existing studies about flight prioritisation mechanisms make use of normative economic models that predict the behaviour of the system under idealised circumstances, such as perfect information and agents' rationality. However, these conditions are often not fulfilled in the real world, where decisions are made in the presence of incomplete or uncertain information, and rationality is limited by the tractability of the decision problem, the cognitive limitations of the decision makers, and the time available to make the decision. In the last twenty years, agent-based computational economics has raised increasing interest as a way to overcome these issues. Agent-based modelling (ABM) allows the observation of the emergent behaviour arising from agents' interactions in a bottom-up process, combining formality and rigour with the minimisation of disadvantages such as strong hypothesis dependency.

2.1 Overview of agent-based modelling

ABM can be defined as a decentralised, individual-centric approach to model a system composed of interacting, autonomous agents. Agents have behaviours defined by simple rules (main drivers, reactions, memory, states...) and interact with the environment and with other agents, influencing their behaviours. When running the simulation, the global behaviour emerges as a result of the interactions of many individual behaviours, showing patterns, structures, and behaviours that were not explicitly programmed into the model, but arise from the agent interactions. ABM offers a way to model social systems composed of agents that interact with and influence each other, learn from their experiences, and adapt their behaviours.

2.1.1 Structure of an agent-based model

A typical agent-based model has three main elements:

- A set of agents, their attributes and their behaviours.
- A set of agent relationships and methods of interaction.
- The agents' environment.

2.1.1.1 Agents

Agents, which can represent people, companies, projects, assets, vehicles, animals, etc., are discrete entities with their own goals and behaviours. The essential characteristics of agents are:

- An agent is a self-contained, modular, and uniquely identifiable entity.
- An agent is autonomous, i.e., it can function independently in its environment and in its interactions with other agents.
- An agent has a state that varies over time.
- An agent is social, i.e., it has dynamic interactions with other agents that influence its behaviour.

Agents may also have other useful characteristics:

- An agent may be adaptive, by having rules or more abstract mechanisms that modify its behaviour. For example, an agent may have the ability to learn and adapt its behaviour based on its accumulated experiences.
- An agent may be goal-directed, having goals to achieve with respect to its behaviours. This allows an agent to compare the outcome of its behaviour with its goals and adjust its responses and behaviour in future interactions.
- Agents may be heterogeneous.

2.1.1.2 Interactions

One of the main elements of ABM is that only local information is available to an agent. Agent-based systems are decentralised systems, which means that there is no central authority that controls agents' behaviour in an effort to optimise system performance. As in real world systems, agents interact with a subset of other agents, called the agent's neighbours. Local information is obtained from interactions with an agent's neighbours and from its local environment. An agent's set of neighbours may change as a simulation unfolds.

The way agents are connected to each other is defined by the topology of the agent-based model, which describes who transfers information to whom. Typical topologies include a spatial grid, a spatial network of nodes (agents) and links (relationships), or a non-spatial model. In some applications, agents can interact according to multiple topologies.

2.1.1.3 Environment

Agents interact with their environment. The environment may simply be used to provide information on the spatial location of an agent relative to other agents or it may provide a richer set of geographic information, as in a Geographic Information System (GIS). An agent's location included as a dynamic attribute is sometimes needed to track agents as they move across a landscape, contend for space, acquire resources, and encounter other situations.

2.1.2 Added value

Most computational modelling research describes systems in equilibrium or as moving between equilibria. Agent-based modelling, however, using simple rules, can result in different sorts of complex behaviour and allows the examination of how the system behaves out of equilibrium, when it is not in a steady state [4][5][6]. A more detailed discussion of the benefits of ABM over other modelling techniques is included thereafter.

2.1.2.1 ABM captures emergent phenomena

Emergent phenomena result from the interactions of individual agents. By definition, they cannot be reduced to the system's parts: the whole exhibits new properties not observed in the individual elements. Classical examples of these phenomena are the behaviour of bird flocks, fish schools, traffic jams, etc.

An emergent phenomenon may have counterintuitive properties that are decoupled from the properties of the parts. As a result, emergent phenomena are difficult to understand and predict. ABM is, by its very nature, the most accurate approach to modelling emergent phenomena: in ABM, the modeller simulates the behaviour of the system's agents and their interactions, capturing emergence from the bottom up.

2.1.2.2 ABM provides a natural description of a system

In many cases, ABM is the most natural modelling approach for describing and simulating a system composed of behavioural entities. ABM makes the model seem closer to reality. As an example, it is more natural to describe how drivers move than to come up with the equations that govern the dynamics of the density of traffic. Because the density equations result from the behaviour of individual vehicles, ABM will also allow aggregate properties to be studied.

2.1.2.3 ABM is flexible

The flexibility of ABM can be observed along multiple dimensions. For example, it is easy to add more agents to an existing model. ABM also provides a natural framework for tuning the complexity of the agents: behaviour, degree of rationality, ability to learn and evolve, and rules of interactions. Another dimension of flexibility is the ability to change the levels of description and aggregation: one can easily play with aggregate agents, subgroups of agents, and single agents, with different levels of description coexisting in a given model.

2.2 Applicability to models of flight prioritisation mechanisms

ABM constitutes a particularly suitable framework to represent and simulate market instruments and other flight prioritisation mechanisms where agent interactions are heterogeneous and can generate network effects. Furthermore, it offers a more intuitive and tractable description of the system units including individual behaviours that would otherwise be very difficult to represent as equations.

Additionally, ABM has prominent synergies with behavioural economics, where deviations from the assumed theoretical behaviour play an outstanding role. The convergence of agent-based modelling and behavioural economics into computational behavioural economics provides a natural framework to incorporate behavioural economics insights about human and institutional behaviour into operational simulation models [7]. Particularly, some of the features that are interesting to analyse are:

- Network effects: the application of certain prioritisation mechanisms to solve the demand-capacity problem in a given time and air sector may cause one or more imbalances in other places later on as a result of the interaction between the agents.
- Different behaviours: flight prioritisation mechanisms are designed so that airlines can decide what action to take when they face air traffic flow management (ATFM) delays. These decisions may respond to different strategies or be influenced by certain cognitive biases which depart from purely rational choices.
- Robustness of mechanisms: assigning different behaviours to airlines allow us to explore the different results that could be obtained with flight prioritisation mechanism and how they would be affected by "irrationalities".

3 Model description

The model used in WP4 to analyse the different flight prioritisation mechanisms, previously defined within WP3, is based on a small-scale agent-based model originally developed in the context of the SESAR Engage KTN Catalyst fund project "Exploring future UDPP concepts through computational behavioural economics".

This model has been extended and improved to align it with the specific objectives of the work package and the overall goals of the BEACON project.

3.1 Overall description

The model simulates *n* days of operations at air traffic level, where the Network Manager takes care of flow management and the airlines make decisions on how to deal with the delays imposed in congestion situations. The model comprises three main elements:

- Geographical context, which provides the environment and the network characteristics for the agents to operate in.
- Exogenous variables, which represent external conditions that affect the model but are not affected by it. They include fuel prices, air navigation charges price, and airlines' cost index.
- Agents. Two types of agents, representing the main actors of the simulation, are considered: the Network Manager and the airlines.

The simulation of one day of operations comprises four main stages:

- In the first stage, with some time in advance (e.g., 2 hours in advance), the Network Manager estimates the future demand for all the sectors within a given period of time (e.g., 15 minutes). This expected demand is checked against the corresponding declared capacity, i.e., the number of flight entries allowed inside the sectors during the mentioned period of time. If the Network Manager detects an imbalance between demand and capacity in a certain sector or group of sectors, it will initiate a regulation and the excess demand will be displaced over time.
- In the second stage, the flights involved in the regulation are provisionally delayed at the origin airport and assigned new take-off times through ATFM slots. The Network Manager calculates these delays and generates the initial provisional sequence following the First Planned First Served (FPFS) principle. This resolution paradigm sequences the flights in the order in which they would have arrived at the constrained airport or sector according to the information present in the filed flight plans.
- The third stage comprises the decision process of the airlines. Once the affected flights receive an initial ATFM slot, the airlines evaluate all possible actions available with the objective of reducing the cost of delay associated with all their affected flights within the hotspot. The number of actions and the complexity of these are defined by the rationale and the flexibility

of the various flight prioritisation mechanisms simulated (e.g., ISTOP, auction and flexible credits[1]). While the ISTOP mechanism builds on the previously calculated FPFS delays, both the auction and the flexible credits generate a totally new flight sequence.

- Finally, the fourth and last stage covers the Network Manager process of acceptance and execution of the prioritisation requests sent by the airlines according to compliance with the traffic restrictions imposed by the model. Once this process is completed, the delays for that hotspot are definitive and the airlines can update the flight plans of their affected flights accordingly.

The first stage is repeated iteratively for each of the time windows into which the simulation day time is divided. Whenever an imbalance is detected, the second, third and fourth stage are performed. When the last time window of the simulation day is reached, the state is saved and the first window of the next day is started. The simulation finishes when the temporal horizon is reached.

3.2 Main assumptions and model restrictions

When approximating a real system into a model, simplifications of the existing phenomena are required until the desired (or feasible) level of abstraction is reached. Moreover, we should add that "if the complexity of the model is reaching a level where we are no longer able to understand the processes involved, the experiments conducted are of little interest and we cannot understand these artificial complex systems any better than we understand the real ones" [7].

Therefore, the model used here is based on some fundamental hypothesis which have been adapted according to the different model extensions and upgrades made in the context of BEACON.

- Flight traffic is checked by the Network Manager in time windows of 15 min and with 2 hours in advance. These values are configurable within a practical range.
- Only ground delays are modelled.
- The demand-capacity relation is based on entries to a sector during a time window. A regulation will be issued every time the entry demand (number of flights entering) exceeds the entry capacity (allowed entries) on a particular sector during a determined time window.
- An ATFM slot can only be left empty if there is no flight that can be located there due to schedule restrictions: no flight can depart earlier than its Schedule Off-Block Time (SOBT).
- If a flight assigned to an ATFM slot creates an additional demand-capacity problem in a different sector and in an already resolved time window, this data will be saved for informational purposes but the slot will be filled anyway. It is assumed that this new imbalance will be solved by other ATFM measures, in the event that it does not automatically disappear due to the inherent variability of the demand caused by future hotspots.
- Airports are approximated as Terminal Manoeuvring Areas (TMA). For the sake of simplicity, the taxi and the runway time are not taken into account, therefore the Actual Off-Block Time

[1] These mechanisms are described in 3.3.4.

(AOBT) is equal to the Actual Take Off Time and in the same way the Actual In-Block Time (AIBT) is equal to the Actual Landing Time.

- Airport curfew is modelled. A closing time is defined for each airport, simulating flight night restrictions where airlines are not allowed to take-off, taxi or land.
- The network topology is built in such way to provide three different disjunctive routes, without repeating any en-route sector, between each origin-destination pair.
- The aircraft speed is constant throughout the entire route.
- Flight cancellations are only considered due to airport curfew.
- No accumulated delays from the previous day are considered.
- Initial delays are randomly imposed on some flights. These delays mimic the possibility of technical failure delays on aircraft.
- Only two different aircraft are included in the model (narrow-body A320 and wide-body B787).

3.3 Main model components

The main components building this agent-based model are the following:

- A **network configuration** that allows the location and tracking of each flight in the model, as well as a definition and structuring of the air space that allows controlling the relationship between declared capacity and expected demand.
- The definition of the **flight variables**. Here we include both the flight schedule, required to provide all the necessary information for the Network Manager to perform ATFM functions, and the passenger connectivity, required to evaluate the impact of each of the different prioritisation mechanisms on the passenger-centred metrics.
- The **agents** of the simulation that represent the actors involved in the appearance and resolution of ATFM conflicts, their roles and behaviour: the Network Manager and the airlines.
- The **prioritization mechanisms** that provide the specific framework of rules and processes on which the airlines rely to carry out their prioritization decisions.

3.3.1 Network configuration

A solid network definition is essential to provide the required information on the spatial location of an agent relative to other agents and to provide a fundamental set of network characteristics. In this model, the network configuration is based on the definition of several airports, sectors and routes.

3.3.1.1 Airport configuration

The defined network consists of 5 different airports, a mix of hubs and secondary airports. **Error! Reference source not found.** shows a more detailed description of the type and characteristics of each one of the airports included.

Table 1. Airport configuration

Airport id	Airport type	Connected with …
Airport A	Regional Airport	Sector 1, Sector 3, Sector 9
Airport B	Regional Airport	Sector 2, Sector 5, Sector 9
Airport C	Hub Airport	Sector 1, Sector 3, Sector 4, Sector 6, Sector 7
Airport D	Hub Airport	Sector 2, Sector 4, Sector 5, Sector 7, Sector 8
Airport E	Regional Airport	Sector 6, Sector 7, Sector 8

3.3.1.2 Sector configuration

The process of sector definition comprises the virtual division of airspace. Thus, the provision of air traffic services is decomposed, in the different sectors, into tasks with manageable workload. The model presented here shows two different types of sectors.:

- En-route sector: nine en-route sectors are modelled, defining different airspace structures crossed by the flights after the departure and before landing.
- Terminal Manoeuvring Area (TMA): one extra sector is defined around each airport simulating a designated area of controlled airspace surrounding an airport, dedicated to take-off and landing operations.

An illustration of the resultant network topology is shown in Figure 1.The white circles designate the id of the particular sector, while the red dots show the connection entry and exit points between sectors.

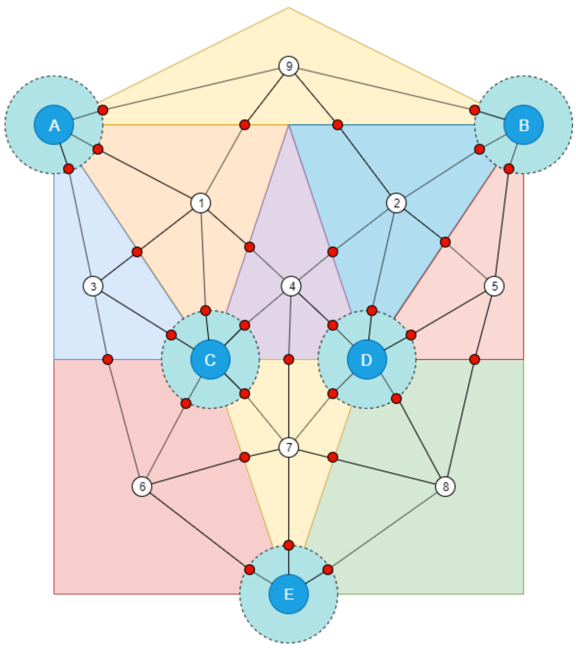

Figure 1. Network topology

3.3.1.3 Route configuration

The route configuration defined for the model follows a fixed trajectory approach with defined entry and exit points for each sector. Additionally, the network configuration is built in such way to allow 3 possible different disjunctive route trajectories for every OD pair. **Error! Reference source not found.** s hows all the different route combinations for the network.

Table 2. Route configuration

OD Pair	Route 1	Route 2	Route 3
Airport A - Airport D (AD)	Airport A - Sector 1 - Sector 4 -Airport D	Airport A - Sector 9 - Sector 2 - Airport D	Airport A - Sector 3 - Sector 6 - Sector 7 - Airport D
Airport A - Airport E (AE)	Airport A - Sector 3 - Sector 6 - Airport E	Airport A - Sector 1 - Sector 4 - Sector 7 - Airport E	Airport A - Sector 9 - Sector 2 - Sector 5 - Sector 8 - Airport E
Airport B - Airport C (BC)	Airport B - Sector 2 - Sector 4 - Airport C	Airport B - Sector 9 - Sector 1 - Airport C	Airport B - Sector 5 - Sector 8 - Sector 7 - Airport C
Airport B - Airport D (BD)	Airport B - Sector 2 - Airport D	Airport B - Sector 5 - Airport D	Airport B - Sector 9 - Sector 1 - Sector 4 - Airport D
Airport B - Airport E (BE)	Airport B - Sector 5 - Sector 8 - Airport E	Airport B - Sector 2 - Sector 4 - Sector 7 - Airport E	Airport B - Sector 9 - Sector 1 - Sector 3 - Sector 6 - Airport E
Airport C - Airport D (CD)	Airport C - Sector 4 - Airport D	Airport C - Sector 7 - Airport D	Airport C - Sector 1 - Sector 9 - Sector 2 - Airport D
Airport C - Airport E (CE)	Airport C - Sector 7 - Airport E	Airport C - Sector 6 - Airport E	Airport C - Sector 4 - Sector 2 - Sector 5 - Sector 8 - Airport E
Airport D - Airport E (DE)	Airport D - Sector 7 - Airport E	Airport D - Sector 8 - Airport E	Airport D - Sector 4 - Sector 1 - Sector 3 - Sector 6 - Airport E
Airport A - Airport D (AD)	Airport A - Sector 1 - Sector 4 - Airport D	Airport A - Sector 9 - Sector 2 - Airport D	Airport A - Sector 3 - Sector 6 - Sector 7 - Airport D
Airport A - Airport E (AE)	Airport A - Sector 3 - Sector 6 - Airport E	Airport A - Sector 1 - Sector 4 - Sector 7 - Airport E	Airport A - Sector 9 - Sector 2 - Sector 5 - Sector 8 - Airport E

3.3.1.4 Network calibration

The topological representation of the network needs to be translated into a physical format. First, each of the lines connecting the nodes in the topology diagram (route trajectory) needs to be assigned a distance. Additionally, air navigation charges need to be modelled in order to cover the air navigation services provided by the Air Navigation Service Providers (ANSPs) over a portion of airspace, in our case coincident with the defined sectors.

According to EUROCONTROL, the charges for the use of air traffic services are computed according to the three following factors:

- Distance factor: the distance factor by charging zone is obtained by dividing, by one hundred (100), the number of kilometres in the great circle distance between the aerodrome of departure (or entry point of the charging zone) and the aerodrome of arrival (or exit point of the charging zone) [8]. For simplicity, the distance values in the model correspond to the great circle distance[2].

- Aircraft Weight Factor: the weight factor (expressed to two decimals) is determined by dividing, by fifty (50), the Maximum Take-Off Weight (MTOW) of the aircraft (in metric tonnes, to one decimal) and subsequently taking the square root of the result rounded to the second decimal [8].
- Unit Rate Factor: the unit rate of charge is the charge in euro applied by a charging zone to a flight operated by an aircraft of 50 metric tons (weight factor of 1.00) and for a distance factor of 1.00 [8].

With the ultimate objective of getting realistic values of cost and distances for each of the routes, each airport was considered to be a representation of a real airport in the ECAC area. In particular, we considered LFBD, EDDK, LFPG, EDDF and LIPE. Consequently, the unit rate factor of each charging zone (which in this case corresponds to each sector) and the route distances can be approximated to realistic values.

Finally, the model is calibrated with values such that the 3 different routes between each OD pair are not equal in terms of cost and distance, neither do they present large differences.

3.3.1.5 Network capacity definition

The declared capacity of each of the sectors defining the network of the simulation needs to be modelled. In reality, the sector capacity depends on a complex combination of factors such as traffic flow direction, coordination procedures, in-sector flight times, etc. For the sake of simplicity, the capacity estimation in the model is only based on the expected entry demand. Consequently, the capacity in the simulation is understood as the number of flights that the Network Manager can safely allow to enter inside a sector during a defined period of time (e.g., 15 minutes).

Given a flight schedule (see section 3.3.2.1 for description) previously generated, the expected entry demand values per sector and time window are computed. With that information the number of allowed entries to each sector during each fixed time window in which the day of operations is divided are generated following a sliding windows approach: the entry capacity of a sector during a certain time window is equal to the maximum expected entry demand for the next 5-time windows. In the event that the maximum expected entry demand is less than 3 flights for a particular sector and time window, that capacity value will be directly assigned to 3 flight entries.

Once the file is generated following all the previous indications, the user is able to manually change capacity values to simulate capacity shortages and manually generate hotspots.

3.3.2 Flight variables

The flight variables are an essential aspect in the model. The flight schedule provides the necessary information for the Network Manager to be able to calculate the expected demand, which is essential

[2] According to [8], for each take-off and for each landing on the territory of a State, 20km are deducted from the total distance for that State. However, due to the approximated network designed for the experiments this is not done for simplicity.

for the air traffic control task. On the other hand, the passenger connectivity is essential for airlines to estimate the cost of delay of their affected flights, which is fundamental for taking proper flight prioritisation decisions.

3.3.2.1 Flight schedule

The flight schedule is one of the most important and critical inputs of the simulation. It includes all the flights involved in the simulation and provides the necessary information about the origin and destination of the flight, departure time, a flight code, the operating airline, the type of aircraft used and an aircraft identifier.

The process of generating this data becomes of great importance to create the right conditions which enable to achieve the objectives of the simulation. This process can be divided in three stages:

Extract and clean real data

In this stage, to recreate a realistic level of traffic in our model, flight schedules of a subset of airports are reconstructed from the information contained in the s06 file[3] of a random traffic day. As the final flight schedule is going to be a combination of real and synthetic data, a typical day of the year with a good volume of traffic is sufficient. Based on that, the chosen day corresponds to May 23, 2019. The process followed to extract and clean the data is explained below:

1. Data filtering:
    a. Repeated entries are deleted.
    b. Flights between the most congested airports in Europe (EGLL, EHAM, LFPG, EDDF, LTFM) are selected.
    c. Real airport ICAO codes are replaced by pseudo codes.
    d. The format of the departure times is changed to the datetime format HH:MM:SS.
    e. Only flights departing from 9:00 to 23:00 are collected.
2. The resulting data are allocated to airlines. To maintain a certain level of realism, and decrease the number of airlines to the desired one for the model, the real airlines are grouped by alliances creating fictitious airlines.
3. The schedule is ordered by departure time and a new column indicating a flight identifier is added.
4. The original aircraft type associated with each flight is changed to one of the two different aircraft included in the model (narrow-body A320 and wide-body B787).
5. A new column is added to the schedule indicating the estimated arrival time of each flight, based on the simulation network.
6. The aircraft ids are generated and assigned to each corresponding flight following a rotation pattern. The approach to generate the rotations works as follows:
    a. Define an aircraft id.
    b. Assign an aircraft id to the first flight in the list of flights.
    c. Calculate the flight arrival time to destination.

[3] Characteristic file format of Eurocontrol DDR2 data source.

d. Check for possible rotations by checking flights departing from the destination airport of that flight and belonging to the same airline. The departure time limit for the rotation will be equal to the arrival time of the previous flight plus the defined turnaround time of the aircraft.
e. Assign the defined aircraft id to the first flight meeting the previous condition. The buffer time is defined as the time between the departure time limit and the schedule departure time of the flight finally chosen as rotation.

This procedure is followed iteratively until there are no flights that meet the conditions, and the process is repeated with the next flight of the list.

The resulting flight schedule contains the following information: origin, destination, departure time, arrival time, airline id, aircraft type, aircraft id, and flight id.

Create point-to-point schedules

The schedule coming from the processed s06 file does not yet meet the desired requirements for the simulation. Almost every airline that was captured in the file corresponds to hub-and-spoke carriers. However, the point-to-point strategy, characteristic of low-cost airlines is still missing due to the previous data cleaning. To include this strategy new flights are included in the schedule. This process is explained below:

1. New "low cost" airlines are defined and a dummy schedule with point-to-point flights between different combinations of airport pairs is generated. The method creates a first flight going from one airport to the other and calculates the flight duration between these two airports. Then, a return flight is generated at the destination airport. The definition of the schedule departure time of the new fight is calculated from the addition of a defined waiting time (turnaround time plus buffer time) to the scheduled arrival time of the first flight.
2. The schedule is ordered by departure time and a new column indicating a flight identifier is added.
3. A new column is added to the schedule indicating the estimated arrival time of each flight.
4. The aircraft ids are generated and assigned to each corresponding flight following a rotation pattern in the same way as in the previous stage.

Merge the schedules

Finally, both the flight schedule coming from the filtered real data and the schedule coming from the generated point-to -point flights are merged and ordered by schedule departure time.

3.3.2.2 Passenger connectivity

To measure the impact that different prioritisation mechanisms have on passengers, it is necessary to include passenger connectivity in the model. Being consistent with the synthetic approach followed for the flight schedule generation, another configuration file was artificially produced with all the information regarding passenger connectivity. For this, the following assumptions were set:

- Only flag carrier airlines have connections.
- The passengers can only have a maximum of one connection in their journey.

- The waiting time for passengers connecting flights lies between 45 and 120 minutes.
- The connections are only between flights operated by the same airline.
- All flights departing from 18:00 onwards are direct flights.
- The number of total passengers inside a particular flight and aircraft is randomised between the following percentages according to airline type:
    - 80-85% of aircraft capacity for hub-and-spoke carriers
    - 85-90% of aircraft capacity for low-cost airlines
- The number of connecting passengers in a flight is computed as the 20% of the total number of passengers of that flight who have not made a connection yet (note the restriction of passengers only having one connection in their journey).
- In the event that the connecting passengers inside a flight could take different second flights, the number of passengers going to each one of these next flights is randomised from the total number of connecting passengers inside the flight.

According to the previous assumptions, the following approach was implemented to generate the connectivity file:

- The total number of passengers is calculated from the total capacity of the particular aircraft and the (random) load factor associated with the airline type.
- The total number of passengers is divided in two different groups: passengers coming from a previous flight and passengers waiting at the origin airport, whose first flight is the current one.
- The possible future flights receiving connecting passengers from the actual flight are computed from the actual flight information, the schedule and the maximum number of connections allowed. These flights are found by filtering the schedule searching for flights operated by the same airline that depart from the actual flight destination airport in a time range going from 45 min to 2 hours, as stated in the assumptions.
- The encountered connecting flights and the final number of connecting passengers boarding them, if any, are saved as connection information of the actual flight.

The generated file consists of several rows corresponding to each of the flights included in the simulation. Each row contains the following flight information:

- Flight id: the id of the flight.
- Pax at origin: number of passengers at the origin airport. These are the passengers that come directly from the airport and not from another flight.
- Pax from connections: number of passengers coming from another flight or flights.
- Pax total: total number of passengers inside the flight.
- Connection info: list of flight ids and number of connecting passengers boarding future flights.

The generation of the passenger objects will follow a 1 to 1 ratio, meaning that a passenger in the connectivity file equals a passenger object in the simulation. Once the passenger objects are generated, they are added to their corresponding airport of origin.

During the simulation, each time a flight departs, the passengers belonging to that flight are moved from the airport to the flight. On the opposite side, when a flight reaches its destination, all the passengers are removed from the simulation, with the exception of passengers connecting with another flight, which are moved to that airport to wait there for their next flight.

3.3.3 Agents

One of the main components of the model are the agents, which, in this case, represent the actors involved in the appearance and resolution of the ATFM conflicts, their roles and behaviour. The main agents identified for this model are: the Network Manager and the airlines.

3.3.3.1 Network Manager

The role of the Network Manager is to apply the corresponding ATFM processes throughout the simulation. It is in charge of the detection of possible demand-capacity imbalances in the air traffic network, as well as of the correct application of the prioritisation mechanisms. There is only one instance of the Network Manager in a simulation.

3.3.3.2 Airlines

The airline agents are the main actors of the simulation. They make decisions to achieve their objectives according to their internal parameters and the environment. They are modelled as cost-minimisers but their final behaviour can be modified by the inclusion of different behavioural biases which depart from purely rational choices. Additionally, their decision-making process is deeply influenced by the different prioritization mechanisms included in the model (see section 3.3.4).

The number of airlines included in the agent-based model is a configurable parameter that depends on the characteristics of the scenarios to be simulated. For these experiments, we only consider 5 airlines classified in two groups, which are differentiated according to their network configuration model:

- Airline 1, Airline 2 and Airline 3: fictitious flag carrier airlines with a hub-and-spoke network configuration.
- Airline 4 and Airline 5: fictitious low-cost airlines with a point-to-point network configuration.

Airline costs are impacted by air traffic control (ATC) charges, the cost of fuel and specially the cost of delay. The calculation of the cost of delay is of special interest for the model because its inherent non-linearity could trigger the use of the available prioritisation mechanisms.

3.3.3.2.1 Airlines cost of delay calculation

The seed of the airlines' decisions within the simulation is the calculation of the cost of delay associated with the different flights affected by demand-capacity imbalances. Two reference documents have been consulted to compute this cost. Regulation (EC) No 261/2004 [9] was comprehensively reviewed to assess the current common rules on compensation and assistance to passengers in the event of denied boarding and of cancellation. On the other hand, the report "European airline delay cost reference values" [10], developed by the University of Westminster and updated in the context of this project in D3.2 [11], was also considered to extract updated reference values for the cost of delay for European airlines. The modelling of these costs in the simulation is based on the following assumptions:

- All the delays are issued when the aircraft is still on the ground, meaning that we are only interested in the at-gate delay cost reference values. For simplicity, we are not modelling any turnaround process.
- Both the maintenance and the crew costs are extracted from the corresponding tables in the University of Westminster's report. The exact crew and maintenance costs per minute are selected based on the base scenario.
- The passenger soft costs are a component in the economics of airline unpunctuality. These are the costs associated with a revenue loss or market value decrease. Unlike crew and maintenance costs, these delay costs are not linear over time. Large delay values are associated with a high-cost value per minute, per person. These costs have been included maintaining their non-linear nature using the appropriate tables found in the University of Westminster's document.
- The passenger hard costs are due to such factors as passenger rebooking, compensation and duty of care. The modelling of these costs is based on Regulation (EC) No 261/2004 and the Articles 91(1) and 100(2) of the Treaty on the Functioning of the European Union (TFEU).
    - Flight cancellations are only included in the model if they are due to the breach of curfew. In that case, it is considered that the final departure of that flight is scheduled the next day.
    - Delays < 3 hours: the airline is required to offer assistance in the form of meals, refreshments and free of charge two telephone calls, telex or fax messages, or e-mails. However, even in the worst-case scenario, this cost is not significant enough for the project so it is not included in the calculation.
    - Delays > 3 hours: In the event of a delay longer than three hours, passengers should be offered a compensation as in the event of cancellation:
        - EUR 250 for all flights of 1,500 kilometres or less;
        - EUR 400 for all intra-Community flights of more than 1,500 kilometres, and for all other flights between 1,500 and 3,500 kilometres;
        - EUR 600 for all flights not falling under (1) or (2).
    - To trigger passenger compensations, we study the passenger arrival delay of both the direct and the connecting passengers. The calculation of the arrival delay of direct passengers is straightforward. However, the arrival time of the connecting passengers extremely depends on the connection or other possible re-booking connections. No extra compensation for loss of connection has been modelled. In the event that the connecting passengers miss their next flight, the airline will try to re-book them in other flights and only in the case that they reach their final destination with more than three hours of delay they will receive the appropriate compensation.
    - There are certain cases in which the airline will also have to pay passengers the cost of accommodation and transportation from and to the airport. Based on literature, this cost is approximated to 150 EUR per person.

- No re-booking: in the event that connecting passengers cannot be re-booked to another subsequent flight, they must spend the night at the airport and therefore require accommodation and transportation.
- Curfew missed: if an aircraft cannot depart due to curfew reasons, either at the destination airport or at the origin airport, the airline must pay accommodation and transportation expenses for some of the passengers, according to the following criteria:
  - All the passengers coming from a connection receive the overnight compensation.
  - Only 50% of the passengers not coming from connections will receive the overnight compensation. This 50% ratio is an approximation for the passengers who are at the home airport and the ones which are not.

3.3.3.2.2 Airlines cost of rerouting calculation

When airlines are notified of an ATFM delay for any of their flights, the developed model includes the possibility for the airlines to execute a rerouting process as an alternative option to participating in the prioritization process. Although this possibility has been disabled for the current version because it is outside the scope of the investigation, with the intention of covering all the possibilities of the model, the calculation of the cost of rerouting by the airlines is briefly detailed here.

The calculation of the cost of changing from one route to another is based on the direct cost associated with flying a certain route. This cost depends on several factors, both external and internal:

- External factors: the total distance of the route and the cost of the air navigation charges.
- Internal factors: the cost index of the airline and the cost of delay function associated with each flight.

The cost index, which corresponds to the ratio between the unit cost of time and the unit cost of fuel, is used to calculate the cost of time for an airline. Consequently, the direct cost of flying a certain route is calculated using Eq1, where fuel burn is the kg of fuel burn per minute; fuel cost is the cost of 1 kg of fuel in Euros; flight time is the duration of the flight in minutes; and the charges cost represent the total cost paid in air navigation charges for flying that specific route.

$$Direct_{cost} = fuel_{burn} * fuel_{cost} + flight_{time} * time_{cost} + charges\_cost \qquad [1]$$

When deciding on rerouting, it is necessary to calculate the difference in cost of flying one route or another. However, to make a fair comparison between the different route options, we must take into account the cost of the delay associated with all alternatives.

This way, the cost of flying the original route is equal to the addition of the direct cost plus the cost of any delay assigned to the flight following that specific trajectory. Likewise, the calculation of the alternative route follows the computation of the new direct cost for that new trajectory plus any possible delay, for instance a potential flight arrival delay due to the new route length.

3.3.4 Prioritisation mechanisms

One of the main objectives of BEACON is to evaluate the performance of different flight prioritisation mechanisms in situations of imbalance between demand and capacity.

To be able to measure this through simulation models, it is essential to implement these prioritisation concepts within the model. This entails a double challenge, because it is not only necessary to adapt the set of rules on which the theoretical framework of the flight prioritisation mechanism is based, but it is also necessary to model the decision process that airlines follow to use this mechanism.

The modelling of this decision process of the agents requires in most cases simplifications that approximate a complex behaviour to a series of equations or heuristics that can be transferred to a computational model.

In this model, each prioritization mechanism is isolated in a component (or class) where the different steps that build the prioritisation process are executed. From within these components, the pertinent calls are made to the agents involved, whose possible actions depend on the mechanism activated at each moment. The output of this process is final sequencing of the flights in the ATFM regulation.

3.3.4.1 Inter Airline Slot Swap Offer Provider (ISTOP)

This mechanism is based on the algorithm developed by the University of Trieste (UNITS) which aims to extend the possibility of inter-airline slot swaps within the user-defined prioritisation process (UDPP) framework.

The key ideas behind the ISTOP mechanism are the following:

- To use the new sequence obtained with the application of the UDPP features FDR and SFP as a new baseline, of which the detailed description can be found in the deliverable D3.1 [12].
- To exploit inter airline slot swap opportunities to further reduce delay costs for the AUs (in addition to the UDPP).
- To keep the UDPP equity philosophy by ensuring that no negative impact on other airlines is created by the mechanism.

The mechanism aims to encourage airlines to provide the Network Manager with a parametrisation of their flights' cost function, which will be partially obfuscated via a normalisation procedure to keep the real cost models confidential while ensuring that the Network Manager has the possibility to evaluate the conveniency of potential slot swaps for the users.

The cost models approximations defined by the parametrisations given by the airlines, once normalised, are called penalty functions and will be used by the Network Manager to compute a set of inter-airline slot swaps. The swaps, called in this framework offers, are generated by an optimisation algorithm that determines the offers which minimise the overall penalty. UDPP equity is preserved as the model constraints ensure that all offers must be convenient for the users involved, i.e., no airline can be impacted negatively. In addition, the minimisation of the overall "penalty", the approximation for the total cost, introduces an additional equity feature: the cost model normalisation is performed for each airline, rescaling the delay cost values of all airline's flight within the interval [0,1]. This is achieved simply considering for each airline the maximum value assumed by its flights' cost functions as a rescaling factor. This procedure allows to flatten the potential difference of delay cost magnitude within the airlines, avoiding airlines with higher delay costs to be favoured at the expense of the ones with lower delay costs. As highlighted also by D6.1 [13], this choice has a big influence on which airlines are favoured by the mechanisms, and should be discussed with the AUs, since there is no perfect "equity", only agreements that are acceptable to all parties.

In order to be consistent with the UDPP philosophy, the offers selected by the ISTOP optimisation do not directly provide the solution which generates the final sequence. Instead, the offers represent simply a what if scenario: they are sent to the corresponding airlines which have, in principle, the right to accept or refuse them. If an offer is accepted by all the parties involved, the underlying slot swap takes place defining the position of the corresponding flights in the final sequence. If an offer is refused by any of the airlines included, all flights involved in the offer keep the same position obtained with the UDPP.

The development of the algorithms for both the UDPP mechanisms and the ISTOP process have been developed by UNITS. However, to be able to execute these algorithms in this agent-based model and in the simulation model used in WP5 (Mercury), it was necessary to develop a Python library[4], named *Hotspot*, which acts as an interface between the different platforms.

The version of ISTOP finally included in the *Hotspot* library includes an automatic approximation of cost functions. Indeed, the airline agents only communicate to the Network Manager agent some simple parameters that are then used to build a cost function. These parameters are called margin and jump, with the margin notionally representing the time after which the flight would cost the jump to the airline (the cost is null before that). This effectively represents a step-wise cost function.

In order to estimate the parameters to be sent to the Network Manager, the airline agents use a simple regression procedure, with the Powell algorithm. In this sense, the parameters sent to ISTOP can be seen as the most honest one, because the approximated functions built back by the Network Manager are the best approximations of the real cost functions.

Also, for simplicity, the previously mentioned condition that airlines can accept or reject swaps found by the ISTOP algorithm has been relaxed. In order not to overcomplicate the implementation of this mechanism, it is assumed that the possible slot exchanges found by the algorithm are automatically accepted by the airlines. Note that this assumption relies on the theoretical efficiency of the ISTOP mechanism itself, in which cost functions are approximated first to a step function and then renormalised. The better the approximation is, the more likely a fully rational airline will accept the swap from the Network Manager.

For the agent-based model, the resolution workflow of an ATFM regulation when the ISTOP is used is relatively simple, mainly because the entire decision process of the airlines is found within the UNITS algorithms, since we assume that the airlines act honestly, and is not explicitly coded within the model, contrary to the other mechanisms explained below (flexible credits and auction). Whenever there is an imbalance between capacity and demand in the network, the following sequence of steps is executed:

1. The Network Manager agent issues a regulation and a first slot allocation is generated based on the FPFS principle.
2. The inputs necessary to execute the UNITS algorithms are encoded in the format required by the algorithms. These are the slots of the regulation and the affected flights, which contain

[4] A Python library is a collection of related modules. It contains bundles of code that can be used repeatedly in different programs.

necessary information, such as their estimated time of arrival to the regulated area (ETA) and cost function (or margin and jump to approximate the cost function).

3. The hotspot resolution algorithm is executed by calling the appropriate function within the *Hotspot* library.
4. The results of the resolution of the regulation are decoded and loaded in the agent-based model. The final slot allocation is generated and the sequence is fixed.

As can be observed, the resolution methodology is extensible to be able to use different resolution algorithms. One could select different algorithms available in the *Hotspot* library, such as:

- globaloptimum: computes the allocation that minimises the overall cost of delay.
- nnbound: global optimum with the additional constraint that the new allocation cannot produce negative impact on any airline.
- udpp*:* computes the new slot allocation based on the optimal UDPP prioritisation defined by each airline. The airline knows the actual cost function of all its flights and uses it to carry out prioritisation.
- udpp_approx*:* udpp but the airlines use an approximated cost function, built with a step-wise function like previously.
- udpp_istop: computes udpp algorithm and then it finds the inter-airline possible swaps further reduce the impact on the airlines involved. This is done with the real cost functions, both for udpp and istop.
- udpp_istop_approx: udpp but the airlines use an approximated cost function, built with a step-wise function like previously. This is the implementation of the procedure described there above.

The experimental scenarios defined for this model only include the use of the algorithms related to the ISTOP scheme (udpp_istop and udpp_istop_approx), although in the future it would be interesting to extend the analysis to other algorithms such as globaloptimum or nnbound to compare the performance.

3.3.4.2 Flexible Credits

For the development of the credit-based mechanism, two possible alternatives were initially proposed. The first one was closely aligned with the work on flexible credits performed by the UDPP team within EUROCONTROL and the second one suggested an innovative method where airlines are able to buy and sell priorities for all their flights of the day through the use of credits. After several discussions with the partners, especially with SWISS, it was decided to approach the second alternative. This second option has some notable advantages over the first one, including that the parameters communicated to the Network Manager are very similar to the ones used in ISTOP.

The mechanism designed is based on the concept that all flights on the day of operations receive at the beginning of the day a default operational value of two parameters that approximate their cost function. The first parameter is the margin and refers to the maximum delay of a flight to avoid non-compliance with certain unwanted events (passenger transfers, compensation, crew or aircraft restrictions, curfews) that drastically increase the cost of the delay. The second parameter, which we call jump, refers precisely to that drastic increase in cost that the airline suffers when that operational margin is exceeded. Note that these parameters are very close in meaning to the ones used in ISTOP to approximate the cost functions, as described before. The main difference is that the jump here can be understood as an absolute cost (in euros), as opposed to ISTOP where the jumps are relative to each other, because of the renormalisation procedure.

In this way, each time the Network Manager issues a regulation, the airlines with affected flights have the possibility of modifying the value of the default parameters of each flight to, ideally, adjust them as much as possible to their real values. In order to do that, airlines are endowed with a certain number of virtual credits at the beginning of the simulation. These credits can then be spent by the airlines to modify the parameters of their flights, and are carried throughout the simulation afterwards.

If an airline decides to reduce the margin value of one of its flights, or increase the jump value, it will have to pay credits. This is because this modification seeks to prioritise that particular flight and adds rigidity and pressure to the optimization system of the final sequence of flights. On the contrary, if the margin is increased and the jump is reduced with respect to its default values, the airline will receive credits, in exchange for removing pressure from the system. The following equation shows how the number of credits to pay or earn is calculated.

$$C = p1(-\Delta m) + p2(\Delta j) \quad [2]$$

Where Δm and Δj are the change in minutes on the value of the margin and the change in euros on the value of the jump, respectively. On the other hand, $p1$ and $p2$ are the prices of changing one unit of margin and jump, respectively. Note that these two prices are parameters of the mechanisms and need to be calibrated in general. More precisely, the ratio between the two prices has to be chosen. Note also that these parameters can be different in different simulations, an option which has been explored in WP5.

Once all the airlines have made a decision regarding the possible change of parameters of their flights, the Network Manager rebuilds approximated cost functions based on these parameters. Next, it optimizes the sequence of flights in the regulation for minimum total cost for all airlines.

The new slot allocation resulting from the optimization is final and fixed, which means that no more credits can be exchanged within that regulation once it has been determined. This system increases the likelihood that an airline will have a better situation by spending credits but it does not guarantee it. The mechanism is intended to balance these possible negative effects for an airline in a regulation with a potential improvement of its situation in future hotspots, thus maintaining equity in the medium term.

3.3.4.2.1 Decision-making process for airlines using Flexible Credits

For the model agents to be able to use this mechanism, a decision process must be approximated and simplified with a series of rules or heuristics. Particularly for the airlines, an 'honest' behaviour has been assumed in which they always try to modify the default parameters to the real values as long as the credits they have allow it. This assumption is considered reasonable because a more faithful approximation to the flight cost function considerably increases the probability of receiving a fair allocation in relation to the importance of that flight, given that the final sequence is based on cost-based optimization.

An agent airline using this credit-based mechanism takes the following actions:

- Identify all the airline flights involved in the hotspot and their associated data.
- Compute the required credits to adjust the default parameters of the affected flights to their "real value", i.e., the values that allow to build the step-wise function which approximates the

best the real costs. It may be that the airline earns credits or has to pay them depending on the flight and the default values of the parameters.

- Update the parameters of the flights whose modification makes the airline earn credits and save those credits (update the credit balance of the airline)
- Update the parameters of the flights where the airline must pay credits.
  - If there are enough credits to change all the parameters of this set of flights to their real values:
    - Update the value of the flight parameters to their real values.
    - Subtract the credits used from the balance of credits that the airline has at that time.
  - If there are not enough credits to change all the parameters of this set of flights to their real values:
    - Set the number of credits that will be used to modify the parameters of each affected flight proportionally to the credits that each flight would need to adjust its parameters to the real value.
    - Update each flight parameter according to the available credits for that flight, prioritizing the margin adjustment over the jump adjustment.

As previously mentioned, the credits are carried throughout the simulation. The credits that an airline earns or spends in a regulation are updated in their own credit balance to be used in future hotspots. Also, the number of credits the airline ends the day with is equal to the number of credits it will start the next day with.

For these experiments the following design decisions have been made when configuring the parameters of the credit-based mechanism.

- The airlines start the simulation with 0 credits.
- The default value of the margin is 60 min and the default value of the jump is 50000 €.
- The price of modifying one minute of the margin value ($p1$) is 1 credit and the price of modifying one euro of the jump ($p2$) is 1 credit.

3.3.4.3 Auction

The auction mechanism provides a method of building the final slot allocation as a result of the amount of money airlines are willing to pay to occupy each of the auctioned slots. However, due to the strong opposition from airlines to pay for the ATFM slots, a virtual currency is introduced to replace real money. See also the discussion in D6.1 [13] about the credit/money issue.

The proposed auction-based mechanism uses a primary auction and virtual credits as an artificial currency. These credits are distributed to all airlines at the beginning of the simulation as an initial allocation that can be based on different criteria (e.g., size of operations). The amount of system credits will remain constant throughout the simulation. Consequently, once distributed, these credits are neither created nor destroyed, they will only be transferred from one airline to another as a result of the different auctions.

Every time the Network Manager finds a demand-capacity imbalance in the network, it issues a regulation. In a first step, it finds the affected flights in the hotspot and issues the ATFM slots where to assign each of the flights. Next, instead of ordering those flights according to the FPFS principle, a primary auction is held where each airline bids a certain number of credits to place their flights in the slot sequence. Therefore, an auction will be executed for each regulation issued during the day of operations.

The auction process that has been designed is based on the following hypotheses:

- The currency used in the auction is virtual credits.
- The airline bids for each affected flight and each ATFM slot where that flight could be placed. The result is a sequence of bids for each affected flight.
- No possible combination of bids between the different flights of the same airline must exceed the number of credits available to the airline at that time.
- Bids can be negative (e.g., -15 credits), meaning the airline receives credits if they end up with that slot.
- A zero-sum condition is required for the bids that the airline sends for each of the different slots where a flight could be placed. The reason is to combat possible abusive behaviour since negative bids are allowed (e.g., an airline bids -10000 credits for a slot just to win credits).

	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5	Sum
A1001	50 cred.	25 cred.	10 cred.	-30 cred.	-55 cred.	0

- The number of credits in the system remains constant. When an auction ends, the excess credits (credits paid by airlines minus credits earned by airlines) are distributed proportionally among all airlines.
- Slots are only auctioned where there is competition between flights, that is, where more than one flight can be placed. If due to schedule restrictions only one flight can occupy a given slot, that flight is directly placed in that slot and will not participate in the auction.
- Flights which have already participated in an auction do not participate in another auction if they are affected by another regulation. The slot corresponding to the time that the airline already bought for that flight in the previous regulation is directly blocked and filled.

The approach designed to implement the auction process within the model consists of these five main sequential steps where the two types of agents included in the model play a key role.

1. Flight checking: the Network Manager agent checks the affected flights to see if any flight has already participated in a previous auction and therefore the airline has already paid for a certain time. In that case, the slot corresponding to that time withing this regulation is assigned directly to the flight and is deleted from the auction.
2. Slot competition checking: the Network Manager agent goes through the ATFM slots to check if there is more than one flight that can bid on them. In the event of no competition, the only flight that can be placed there, if any, is assigned directly to that slot and both are removed from the auction.
3. Airlines' decision process: the airline agents decide the slot bidding sequence for each of their affected flight.

4. Auction assignment: the Network Manager agent assigns each ATFM slot to the winning flight based on the all the bids received and generates the final slot allocation.
5. Credits redistribution: the Network Manager agent calculates the excess credits and redistributes them among all participating airlines to ensure that the number of credits in the system remains constant.

3.3.4.3.1 Decision-making process for airlines using Auction

In order to represent how the airline decides the value of the bids, its decision process has been simplified into a series of heuristics that approximate a reasonable behaviour. An agent airline using the auction mechanism takes the following actions:

- Get the affected flights in the hotspot that can participate in the auction.
- Sort the flights according to the average cost of delay in that hotspot, which is the average cost of the flight over all slots in the regulation.
- For each flight X in decreasing order of priority:
  - Calculate the cost of delay associated with each available slot in the auction. If due to time restrictions the flight X cannot be placed in any of those slots, the value of that cost will simply be None.
  - Find all the possible combinations of bids that could be awarded from the sequence of bids already generated for previous flights in this regulation, if any. The more flights affected, the more different combinations there can be, always bearing in mind that the same slot can never be allocated to two different flights, since only one flight can be placed per slot.
  - Compute the maximum number of credits that can be offered for each slot for flight X according to the bids already decided for other flights, should they be awarded, based on the combinations previously computed, if any, and the balance of credits of the airline. The airline can never bid more credits for a slot than it has.
  - Compute the bid for each possible slot as the subtraction between the average cost of delay for the flight in the hotspot and the cost of delay in that particular slot, multiplied by a conversion factor between euros and credits.

$$bid = (cost_{avg} - cost) * ratio_{credit-EUR} \qquad [3]$$

  - Decide the final bid by comparing the value of the calculated bid and the limit set previously:
    - If the calculated bid is higher than the limit, the final bid is equal to that limit.
    - If the calculated bid is lower than the limit, the final bid is equal to that calculated value.

As mentioned above, to carry out the auction process in the simulation, certain design parameters have to be chosen that can potentially influence the performance of the mechanism. It is especially necessary to determine the number of credits with which the airlines start the simulation, since this determines the number of credits in the system. To determine this number, a series of calibration simulations have been carried out with different initial credit values to finally choose the case that optimised performance for the chosen experimental scenarios: this is 700000 credits per airline. However, this value is configurable and the calibration could be refined in a future work.

On the other hand, it is also important to define the equivalence system between credits and euros that airlines need to calculate their bids. For simplicity, the ratio has been assumed to be 1 to 1, although again this is user configurable.

3.3.5 Behavioural models

BEACON is primarily focused on overcoming the typical rigid assumptions and hypotheses made by classical approaches when modelling certain economic mechanisms by using behavioural economics instead.

The proposed approach is based on including some selected behavioural biases within the decision-making of airlines in situations where they face ground delays due to demand capacity imbalances in the airspace system, so we can evaluate the impact of the prioritization mechanisms in a more realistic way.

The behavioural framework chosen for the BEACON project is closely tied to the decision-making of interest within the ATM agents represented in the model and is based on the main pillars of behavioural economics: prospect theory and hyperbolic discounting.

The theoretical background on these behavioural models is detailed in the document D4.1 [14]. Here, we will focus on detailing their implementation within the toy model with a special emphasis on the limitations and simplifications made.

3.3.5.1 Prospect theory

There are three main ideas behind prospect theory:

- Individuals make decisions based on their experienced value of losses and gains relative to a fixed reference point. This is in contrast to standard utility theory where utility is generally calculated based on net wealth.
- Individuals are more sensitive to loss that to gains, e.g., they are twice as unhappy with a loss of X euros than they are happy of the gain of X euros.
- Individuals see their “extra unhappiness” decreases more slowly as a loss increases.

These three ideas are added to the basic utility theory, where the “extra happiness” of an individual increases more slowly as gain increases. As explained in more detail in D4.1 [14], this is easily captured by a “prospect function”. This function has four characteristic mirroring the above characteristic of prospect theory and the one from utility theory:

- The utility function is centred around fixed reference point which in general does not relate to zero wealth.
- The utility function has a higher slope in the negative region than in the positive one around the point of reference.
- The utility function is concave in the loss domain.
- The utility function is convex in the gain domain.

Here we are using this function calibrated with values extracted from literature, as explained in D4.1 [14].

3.3.5.1.1 Prospect theory implementation in ISTOP

As explained in the section 3.3.4.1, the implementation of the ISTOP mechanism in the model is based on the execution of different algorithms that optimize the final sequence of flights given a series of information on the costs of each flight from which the Network Manager extracts the priorities of each airline and finds potential inter-airline swaps.

Given the implementation characteristics of this mechanism in the model, the best way to apply the prospect function is to include it in the inputs which enter the ISTOP algorithm. Being more precise, the prospect function is used to directly modify the delay cost function of each affected flight as a previous step to the execution of the prioritization algorithms, effectively distorting the costs communicated to the Network Manager.

To be able to apply the prospect function to the calculation of the cost of delay, it is required to establish a fixed reference point by which the airline differentiates between losses and gains. After the feedback received from Eurocontrol and SWISS, and numerous consultations with UoW and Salient, the reference point chosen corresponds to the zero-delay position for the flight in the hotspot. In this way, we modify the cost of delay function of each flight so that the cost calculated for each slot refers to the perceived cost with respect to the reference point and not to the real cost.

Sometimes, it is possible that the flight already arrives with an accumulated delay (non-ATFM delay or a previous ATFM delay from a regulation already solved) that must be taken into account for the cost calculation. This means that even if there is a position for that flight in which the delay is zero for that hotspot, the cost of the delay will not be zero because it already has an accumulated delay that must be considered. This point is still the reference point for the airline, illustrating how reference points are independent of actual net wealth for the agent.

3.3.5.1.2 Prospect theory implementation in Flexible Credits

The way in which the flexible credit-based mechanism is implemented in the model does not seem to include the possibility of airlines falling into this type of bias.

As explained in detail in section 3.3.4.2, airlines do not calculate here the cost of the delay to decide on the prioritization of their affected flights, but rather try as far as possible to modify the default values of the margin and the jump to the actual values of the cost function of the flight. In this way it is very difficult to infer what could be the reference point by which the airline differentiates between potential losses and gains which makes it impossible to apply the prospect function.

3.3.5.1.3 Prospect theory implementation in Auction

According to the bidding process described in section 3.3.4.3, the most direct way to incorporate the prospect theory into the airlines' decision process is through the calculation of the cost of delay they perform to generate the final bids.

Analogously to the ISTOP mechanism, it is necessary to establish a fixed reference point by which the airline differentiates between potential losses and gains. The same reference is used: the zero-delay position for the flight in the hotspot.

The choice of this reference point comes with the fact that any situation other than that zero-delay position constitutes a loss. If we see this from the perspective of the prospect theory function curve it means that the only area of importance is the loss domain [14]. Therefore, the 'loss-aversion impact' is lost as there is no gain domain.

The application of the prospect function fits within the decision-making process of the airline, right at the step where the airline decides the value of each bid (see section 3.3.4.3.1). The following example illustrates the process.

1. The airline calculates the cost of delay of its only affected flight (A1001) in the hotspot.

	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5	Sum	Ave.
A1001	0 €	10 €	15 €	25 €	30 €	80 €	16 €

2. The airline computes the loss with respect to the reference point (zero-delay position). For this particular example the flight A1001 has an *eta* corresponding to the first slot in the sequence.

	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5
A1001	0 €	-10 €	-15 €	-25 €	-30 €

3. The airline applies the prospect function defined in [14] to the relative losses calculated in the previous step with $\alpha = \beta = 0.88$ and $\lambda = 2.25$.

	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5	Sum	Ave.
A1001	0 €	17 €	24 €	38 €	44 €	123 €	≈ 25 €

4. The airline follows the same procedure defined in the section 3.3.4.3.1 for the final decision on the value of the bids, with the only difference that the costs are now modified by the prospect theory. Next, the conversion of euros to credits is made, which for simplicity is established from one to one.

	Slot 1	Slot 2	Slot 3	Slot 4	Slot 5
A1001	25-0= **25**	25-17= **8**	25-24= **1**	25-38=**-13**	25-44=**-19**

The resulting bid sequence after this process is different from the sequence calculated with the 'rational' strategy and can potentially generate a different order of flights in the slots. In the illustrated example the losses are relatively low, of the order of tens. This implies that after applying the prospect function, the value function is greater than the real one due to the convex shape of the prospect function (because the initial slope in the loss domain is higher than 1). However, within the model, the values of the cost of the delay and therefore the losses calculated by the airlines are much higher. Consequently, the value of the function calculated by applying the prospect function is less than the real value, due, again, to the concave shape of the function in the loss domain.

3.3.5.2 Hyperbolic discounting

As described in detail in the deliverable D4.1 [14], hyperbolic discounting is a time-inconsistent form of temporal discounting whereby the subjective value within shorter timeframes falls more rapidly, and at points further in the future, these discount rate falls more slowly in longer delay periods.

After several internal discussions, the implementation of the hyperbolic discounting bias to the behaviour of the airlines agents within this model have been approximated only with the inclusion of the 'present bias'. Present bias is the tendency to rather settle for a smaller present reward than to wait for a larger future reward, in a trade-off situation. It describes the trend of overvaluing immediate rewards, while putting less worth in long-term consequences.

The reason for not implementing hyperbolic discounting to its full extent and only limiting the study to this bias is directly related with the time scale of the defined simulation scenarios (see section 4) which is of the order of days. This time scale is not enough to truly appreciate the effect of the hyperbolic discount, whose parameters are calculated in the context of years.

3.3.5.2.1 Present bias implementation in ISTOP

The prioritisation mechanism based on the ISTOP algorithm does not present any situation where the airlines must decide between whether to receive an immediate reward instead of waiting to get another larger reward in the future, since there is no credit involved. In this way, the appropriate conditions are not given for this bias to appear, and consequently it will not be implemented in this mechanism.

3.3.5.2.2 Present bias implementation in Flexible credits

The flexible credit-based mechanism offers an opportunity where airlines might prefer to receive a reward more immediately, for example by spending more credits to improve the margin and jump values of their flights, even above the real ones, instead of waiting to receive a potentially larger reward in the future, for instance by saving more credits that could be valuable to adjust the parameters of other important flights that may be affected in the future.

The approach chosen for the implementation of this bias is based on the application of an increase factor modifying the number of credits that the airline spends in the regulation. Consequently, the jump and margin values are slightly displaced from the honest ones, tighter margin and bigger jump, increasing the chances that the airline ends up in a much more favourable situation.

If this bias is activated in the simulation scenario, it will only affect the airlines in the situations where the airline has enough credits to at least adjust the parameters of its flights to the real values defined by its cost function. The rationale behind this is that you can only spend extra credits to try to end up in a better situation in the current hotspot (immediate reward) if you have extra credits to spend.

After a detailed review of the literature, it was found that an approximate factor of 1.13 is a reasonable value to represent the present bias in the model. However, to observe possible changes in the final sequence of flights due to this bias, the value of this factor has been slightly varied depending on the type of airline. This is a realistic approximation since depending on the characteristics of the airline, low-cost or carrier, it can present a different behaviour. Based on that, the following assumptions have been made for the model:

- Low-cost airlines (Airline 4 and 5) have a higher value of this factor, which in the model translates into a factor of 1.2, this represents a behaviour with a more focused vision in the present
- Carrier hub-and-spoke airlines (Airline 1, 2 and 3) have a lower value that ranges between 1.05 and 1.15, which represents a behaviour with a vision a little more focused on the long term.

3.3.5.2.3 Present bias implementation in Auction

As in the previous mechanism, the auction mechanism also offers the ideal conditions for the present bias to appear. The airline must decide whether to spend more credits to bid for a certain slot or reduce the bid and save more credits for possible future auctions.

Again, the approach chosen for the implementation of this bias within the airline's decision process is based on the application of an increase factor to the credit bid that the airline calculates following the 'rational' strategy of generating offers.

This way, the number of credits that the airline bids increase in absolute value, meaning that the airline pays more credits than in the ‘rational’ scenario. This fact is in line with the idea that the biased airline prefers to slightly outbid to have more chances of winning the best slots at the current hotspot, reducing the associated cost of delay now (immediate reward), than saving credits with a lower bid to use them on future regulations.

The values of the factors applied to each airline are the same as those defined for the flexible credit’s mechanism (see section 3.3.5.2.2).

4 Definition of simulation scenarios

For the testing and evaluation of the different prioritisation mechanisms selected, it is essential to define a set of simulation scenarios in which the mechanisms can be used by the AUs.

As defined in the deliverable D3.1 [12], a scenario is considered here as a *set of values for input variables to model*. The definition of this scenario is mainly based on the following aspects:

- The traffic setup. This category includes all the specific schedules, flights, passengers etc. included in the simulations. They form the backbone of any scenario.
- The traffic conditions. This category gathers all variables that are linked to the conditions in which the traffic is realised, for instance delay levels, which in this case is directly linked to the capacity configuration.
- The active mechanisms. These variables simply indicate the presence of the mechanisms that the project wants to test.
- The active biases. The variables include the behavioural parameters for different agents, particularly the ones driving non-rational behaviours in these agents.

4.1 Traffic setup

The traffic set up includes the definition of all the variables related to flights, passengers and network (airports, sectors, etc).

For reasons of simplicity, only one traffic setup has been defined that remains constant for all simulations. The network used is the one described in 3.3.1 and the flight schedule and passenger itineraries are created based on the processes detailed in 3.3.2.

In order to evaluate the mechanisms over time, the same traffic setup has been repeated a total of 50 times. Assessing a mechanism in a single day of operations can lead, in some cases, to ignoring its true impact on the network, especially if we consider its disaggregated effects by AUs. The fact of repeating the same simulation day *n* times gives us the opportunity to observe how that performance varies over time and offers us a much more reliable image of reality.

4.2 Traffic conditions

Traffic conditions encompass parameters that are not included in the traffic configuration but have a direct impact on it. In the specific case of this model, we can find the following variables:

- ATFM regulations,
- non-ATFM delays,
- taxi times, turnaround times, etc.,
- passenger connecting times.

Among these variables, we chose to fix all of them except for regulations. In this model, ATFM regulations can be created artificially or can appear spontaneously due to network effects. Its appearance depends directly on the relationship between demand and capacity, which is given by the flight schedule and the capacity file (see section 3.3.1.5) respectively.

With the intention of creating different traffic conditions for each of the *n* days of simulation, it was decided to create a set of different capacity files where a first hotspot is artificially generated by manually lowering the capacity of a random sector during one or more time windows between 13:00 and 14:00. This way, as the capacity configuration is changing for each day, different results are observed each time, even if the traffic setup remains constant.

The original idea was to design isolated scenarios where a single ATFM regulation focused on a single big hub, a regional airport and a sector in the central airspace, first independently and then at the same time. However, due to the inherent conditions of the agent-based model (network effects) it has finally been decided not to explicitly force each of these types of regulation. Therefore, for each day a first ATFM regulation is defined by the artificial capacity shortage included in the corresponding capacity file. This first regulation can be of the three previous types indistinctly. Later, due to network capacity-balance effects created by the resolution of that hotspot, other type of regulations can be created elsewhere in the system.

One of the fundamental objectives of the definition of the traffic conditions resides in reproducing the adequate conditions so that the AUs can make use of the prioritisation mechanisms so that we can measure their impact on ATM performance. With the first tests it became clear that some of the mechanisms showed their true potential when the delays of the flights affected by the regulations were large enough to potentially cause big losses due to, for instance, compensation costs for passengers. However, given the scope of the toy model (geographical scale and number of flights), that level of delay is very difficult to achieve with ATFM delays alone.

For this reason, it was decided to include a series of non-ATFM delays that affect a series of flights in the simulation. In this way, when these flights, which already have an initial delay, are impacted by a regulation, the corresponding ATFM delay is added to the one already suffered, facilitating the possibility that this new delay will have a great impact on the airline's costs (e.g., breaking a margin) and thus unleashing the true potential of prioritisation mechanisms.

4.3 Active mechanisms

Since the ultimate intention of the experiments is the evaluation of the flight prioritisation mechanisms, the definition of a scenario is directly influenced by the mechanism that will be available to the airspace users. These mechanisms are described in more detail in deliverable D3.1 [12] and their implementation is detailed in section 3.3.4. They are summarised in the following table

Table 3. List of mechanisms included in the simulations

Mechanism tag	Mechanism
udpp_istop	Inter Airline Slot Swap Offer Provider (ISTOP) with true cost functions
udpp_istop_approx	Inter Airline Slot Swap Offer Provider (ISTOP) with approximated cost functions
credits	Flexible credits
auction	Auction

4.4 Active bias

The last aspect influencing the design of the simulation scenario is the presence of human biases within the model agents. As explained in section 3.3.5, these behaviours are embedded in the decision process of the agents by modifying some strategies or behavioural parameters. For example, an airline agent must decide how many credits to use to try to prioritise one of their important flights that has been affected by a regulation. If the *present bias* within the hyperbolic discounting theory is active, the agent will give more importance to receiving a potential reward in the present. In this way, the expected value of the credits is discounted and the airline ends up spending more credits in the present moment to try to receive the reward at that very moment.

As detailed in section 3.3.5, due to scope limitations and the characteristics of the model and the flight prioritisation mechanisms, not all the bias have been implemented in each prioritisation mechanism included in the study. The following table summarises this.

Table 4. Summary of the bias implemented for each prioritisation mechanism

Mechanism tag	Prospect theory	Hyperbolic discounting
udpp_istop	Implemented	Not implemented
udpp_istop_approx	Implemented	Not implemented
credits	Not implemented	Implemented
auction	Implemented	Implemented

5 Analysis of results

This section presents the results of the simulations carried out for the different scenarios detailed above. The performance of each mechanism has been measured according to the most relevant metrics collected in the BEACON performance framework, previously elaborated within WP3 [12].

The results of the metrics are calculated as the mean of the values calculated over 50 simulation days.

5.1 Operational Efficiency

This KPA was termed Punctuality and Predictability in the first version of the BEACON performance framework. However, to align with the nomenclature of the latest version of the SESAR Performance Framework [15], the name has been changed to Operational Efficiency.

According to the SESAR Performance Framework, this KPA covers elements related with predictability and includes punctuality aspects that are related to temporal quality of service aspects of ATM. However, the final metrics computed here only refer to the on-time performance dimension, leaving out the predictability metrics. This decision can be explained by the limitations in the scope of the simulation model: neither the ability of the airlines to change the cost index (change the flight speed) nor the assignments of en-route delays (e.g., holding patterns) are modelled.

The punctuality metrics chosen are two: the average flight arrival delay, which is the total flight arrival delay in minutes divided by the number of flights in the simulation day; and the average passenger arrival delay, calculated as the total passenger arrival delay to final destination in minutes divided by the number of passengers in the simulation day. Figure 2 and Figure 3 show the results of these metrics for each of the flight prioritization mechanisms selected in the simulation scenarios.

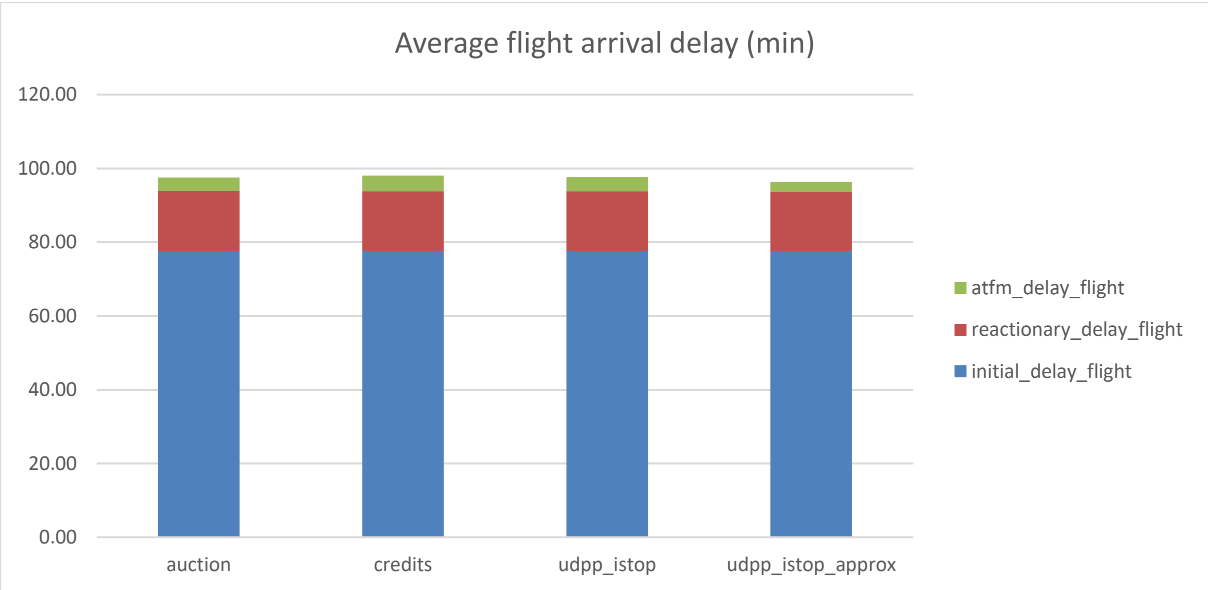

Figure 2. Average flight arrival delay

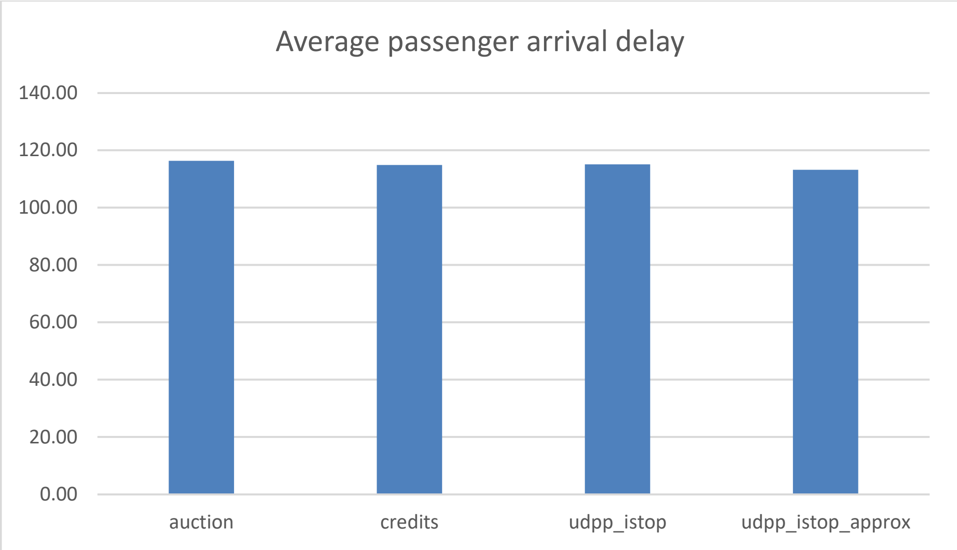

Figure 3. Average passenger arrival delay

Figure 2 shows the average delay per flight for a simulation day for each mechanism, differentiating between three types of delay: the initial delay, which is the artificial delay assigned to each flight at the beginning of the simulation (see section 4.2); the ATFM delay, which is the delay issued to the flights in the regulations due to demand-capacity imbalances; and the reactionary delay; which is the flight delay caused by a late arrival of the aircraft.

All the mechanisms start with the same initial delay conditions. Then, due to the different network effects caused by different slot allocations depending on the mechanism, they show a slight variability in both the ATFM delay and the reactionary one. However, in general, the results show a practically identical behaviour for all the mechanisms, with an average value of delay per flight between approximately 96 and 98 min. This value is very high due to the traffic conditions (initial delays) imposed in the different scenarios. This causes in some cases the cancellation of some flights due to loss of curfew, which also explains the high levels of delay.

Table 5 complements this punctuality analysis with the values of the percentage of flights that are affected by regulations. All the mechanisms present a very similar average with slight variations due to the different network effects.

Table 5. Average ratio of flights involved in regulations w.r.t total number flights

Mechanism	Average ratio of flights involved in regulations w.r.t total number flights
auction	15.00%
credits	14.50%
udpp_istop	13.34%
udpp_istop_approx	11.15%

On the other hand, Figure 3 shows the average delay per passenger for a simulation day. Once again, it can be seen that the values are very high, between 113 and 117 minutes per person, with no significant differences between mechanisms. As expected, the relationship with the metric of the average delay per flight is direct, the later the flights arrive at their destination, the more delayed the passengers accumulate. In fact, the average delay per passenger is even higher due to connecting passengers, who, in the case of missing a connection due to the delay of their first flight, must be rebooked and may arrive at the final airport with much more delay.

According to these results, it is clear that punctuality metrics are not effective in evaluating the performance of the different flight prioritization mechanisms. The results are much more influenced by the initial traffic conditions than by the performance of the mechanisms. This makes sense because the main objective of the flight prioritization mechanism is not to reduce the delay of the airline in the regulation, this is not possible, but to provide the airline with enough flexibility to reorder its affected flights and relocate that delay in its flights in a way that minimises its cost of the delay.

5.2 Cost Efficiency

In the Cost Efficiency KPA collected in the BEACON performance framework we only focus on the cost impact for AUs. This dimension includes both the delay costs and the direct gate-to-gate ANS cost, although for this analysis we will only focus on the delay cost metric. The main reason is that the direct gate-to-gate ANS costs are exactly the same for all scenarios because they depend on factors that remain constant for all cases, for example the flight trajectories flown and the price of fuel.

To analyse and compare the impact of each mechanism on the airlines cost of delay, we have used a metric that computes the ratio between the airlines cost of delay based on the original FPFS slot allocation and the airlines cost of delay based on the allocation using the flight prioritisation mechanism, for one simulation day. This is a clear measure of the mechanism's performance in reducing the cost of delay of the airlines.

Figure 4 shows the average value of this ratio for the 50 days of simulation for each flight prioritisation mechanism. It should be noted that only flights that have been affected by ATFM regulations are taken into account for the metric calculation, regardless of whether they also have an initial and/or a reactionary delay. To calculate the costs of these flights, all types of delays are taken into account.

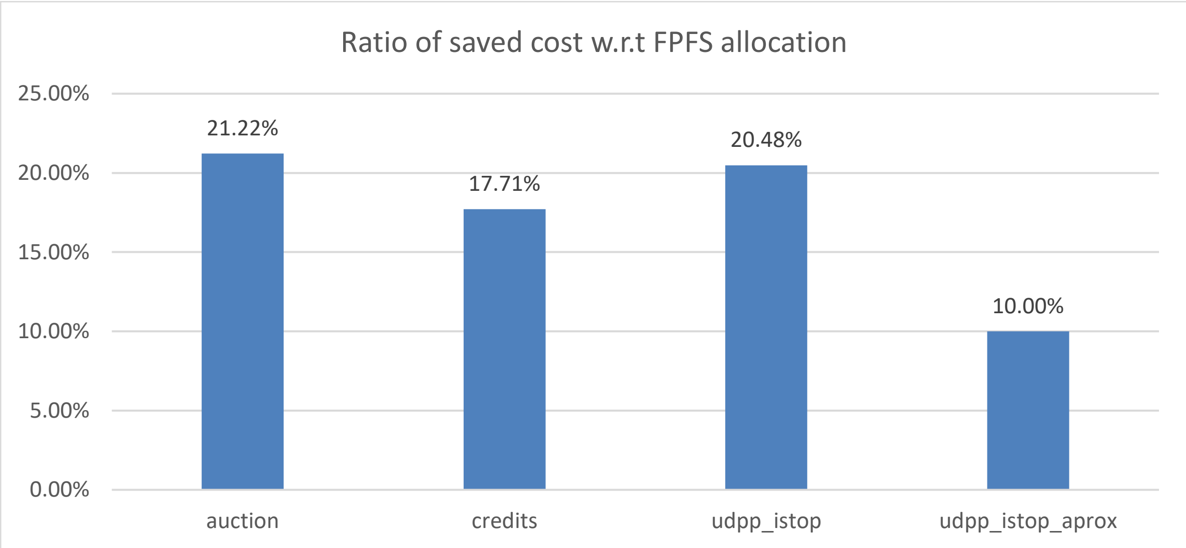

Figure 4. Ratio of saved cost of delay w.r.t FPFS allocation for all the mechanisms

From Figure 4 it can be seen that the mechanism with the highest reduction in the cost of the delay is the auction. The modelled auction proposes a situation where each airline bids for slots according to the importance of the flight based on the cost of the delay. This means that in situations where the credits are sufficient, the final sequence will be totally based on a cost minimisation schema and the allocation will be close to being optimal from that point of view.

Additionally, it is observed that the udpp_istop mechanism also offers a very good performance, close to the auction. This mechanism executes the udpp prioritization process and then computes inter-airline swaps between airlines whenever both benefit from the change. However, since changes are only possible in pairs with a given sequence from the udpp, the reordering of flights is limited and in many cases no swap is possible. In fact, in the preliminary tests run it was clear that ISTOP offers were only found in hotspots with many flights and with very long delays (close to breaking some operational margin). Given the characteristics of the toy model, this situation rarely occurred and for this reason a series of initial delays were applied to the flights that would help artificially create these situations.

The credit mechanism also provides good results, but the performance is slightly lower than the previous two. As explained in section 3.3.4.2, it optimises the flight sequence according to the approximation of the flight cost function which is based on the margin and jump values that the airline can modify with the credits. In this way, if the airline does not have enough credits or the approximation is not good enough, the mechanism performance may worsen. For this particular case, the airlines have enough credits to update these parameters throughout all the simulation days (see Appendix A), so the reduction in performance compared to the auction and udpp_istop comes exclusively from the quality of the cost function approximation.

Finally, it is interesting to see that the udpp_istop_approx algorithm is the one offering the worst results, although it still realises a considerable improvement with respect to the allocation based on FPFS. This mechanism is equal to the udpp_istop with the only exception that the cost functions used by the Network Manager to search for swaps are approximated by a margin and a jump (see section 3.3.4.1). For this case, it seems that these approximations are not good enough, which means that some swaps that would be beneficial do not occur and others that may not be favourable are executed, considerably worsening the performance with respect to udpp_istop.

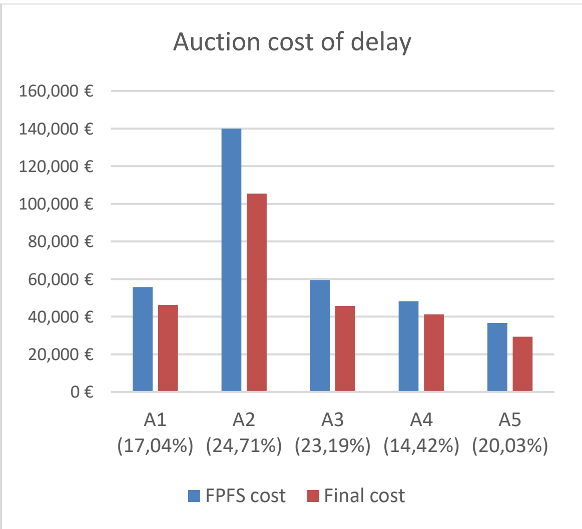

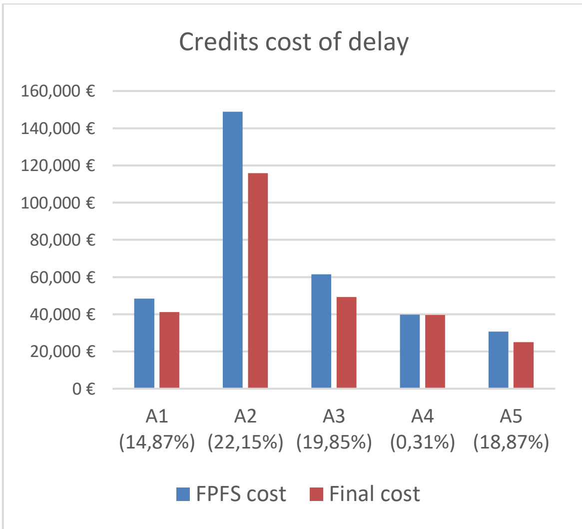

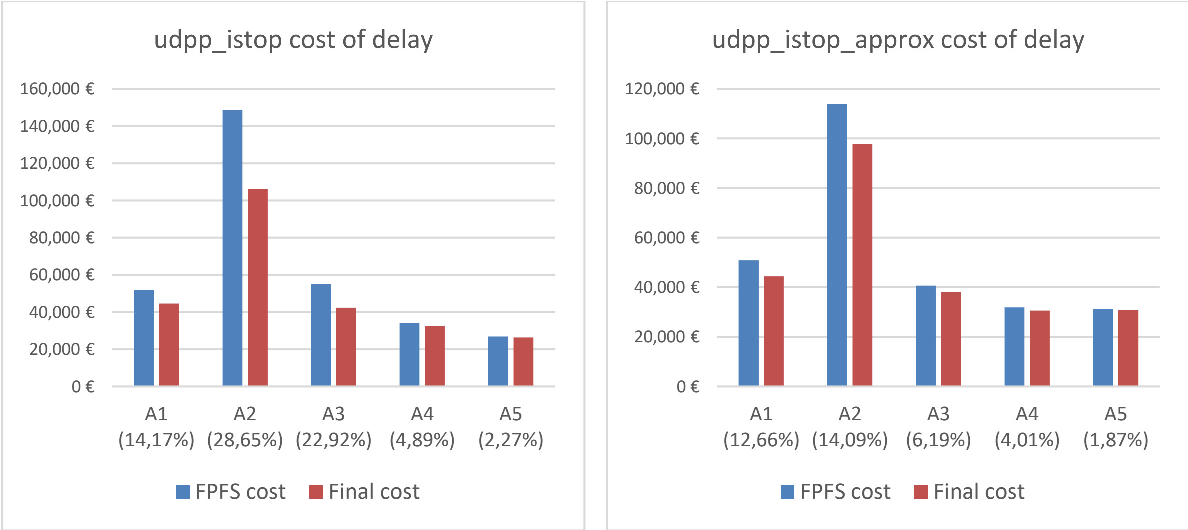

Figure 5. Comparison between costs with FPFS allocation and final costs disaggregated by airline

Figure 5 provides a disaggregated view of the cost of the delay for each mechanism, again differentiating between the costs of the delay based on the FPFS allocation and the final sequence. This perspective is very interesting because it allows us to see how each airline is impacted by these costs and to what extent they improve their situation thanks to flight prioritization mechanisms.

Firstly, it can be seen that the degree each airline is affected by delay costs due to ATFM regulations follows a common trend for each of the four scenarios shown. In fact, three levels can be distinguished, which are deeply influenced by the scenarios design: airline A2 is always the one with the highest cost of delay; airlines A1 and A3 have a similar daily FPFS cost of between 50,000 and 60,000 euros; finally, airlines A4 and A5 have a FPFS cost between 30,000 and 40,000 euros.

Another aspect to highlight from these results is that it is evident that all the airlines improve their situation with respect to the baseline based on the FPFS principle. This is very important because it means that even if an airline is punctually negatively impacted in a regulation, the situation of the airline in the medium term has proven to be positive. This will be discussed in more detail in the next section.

5.3 Equity

Within the UDPP programme the equity is the idea by which the actions of one AU must not negatively affect (significantly) other AUs' flights. This supposes an essential requirement from the airline's perspective and therefore it is essential to evaluate how the different flight prioritization mechanisms respect this constraint.

To assess equity, different metrics have been calculated that are based on the indicators suggested in the SESAR Performance Framework. First of all, Figure 6 shows the change in the costs of each airline with respect to the other airlines. This indicates the percentage of improvement that each airline takes with respect to the total improvement for all airlines together, and the improvement in this case means the reduction in delay costs with respect to the FPFS allocation.

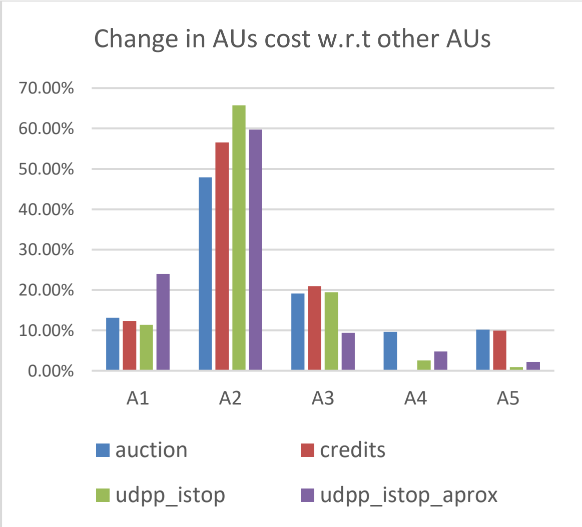

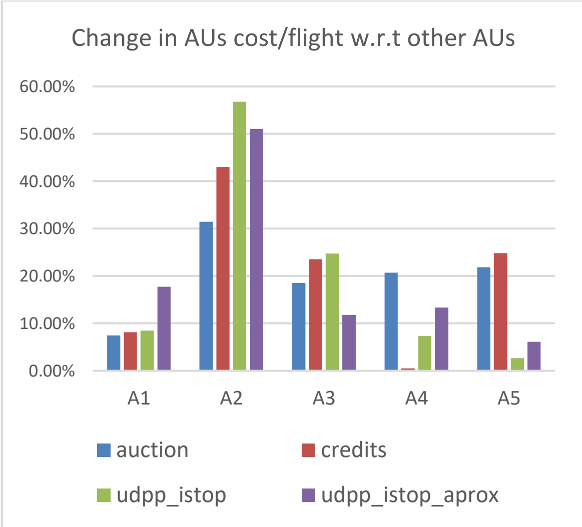

Figure 6. Change in AUs costs w.r.t other AUs in absolute value and relative to the flights of the airline

Figure 6 shows this percentage both in absolute value and per flight where each colour indicates a different prioritization mechanism.

Given these two metrics, the most equitable situation should be one where all airlines take equal advantage of the performance improvement provided by the mechanism. In this case, since there are five airlines, that would mean that they all reach a percentage of 20%. Then, the results show that, a priori, there is no fully optimal equity for any of the mechanisms tested. However, differences can be seen between them.

On the one hand, the auction is the mechanism that offers more favourable results in terms of equity. If we analyse the equity metric per flight we can see that three of the airlines (A3, A4 and A5) are around 20% while airline A2 accounts for more improvement, approximately 30%, to the detriment of airline A1 with only around 10%. With sufficient credits, the auction provides with total flexibility to bid based on the importance of flights and, especially, also guarantees the same possibilities of prioritisation to all airlines regardless of the number of flights affected. This may be the reason behind the good equity results.

At the other extreme is the udpp_istop mechanism. Here the imbalances between airlines are more notable, especially because airline A1 accumulates almost 60% of the reduction in total cost, while airlines A4 and, especially, A5 obtain practically no improvement. This is mainly caused by the fact that none of the flights of these airlines participate in the found swaps, due to the lower number of flights of these low-cost airlines and the traffic conditions given by the flight schedule used.

The results for the credit mechanism are also interesting. In this case, airline A4 ends up in a situation very similar to the baseline (FPFS). Due to the traffic and network configuration, favourable conditions are not generated for this airline to notably reduce the cost associated with its impacted flights: the delays of its flights are rarely close to breaking some operational margin, so its final allocation after the prioritisation process does not represent such a considerable saving for the airline.

Figure 7 shows another perspective of the degree of equity offered by each flight prioritisation mechanism. The metric computed is the percentage of cost saved by each airline with respect to the allocation based on FPFS for each mechanism, or in other words, how much the airline improves its situation by using the prioritisation mechanism compared to the baseline situation (allocation FPFS).

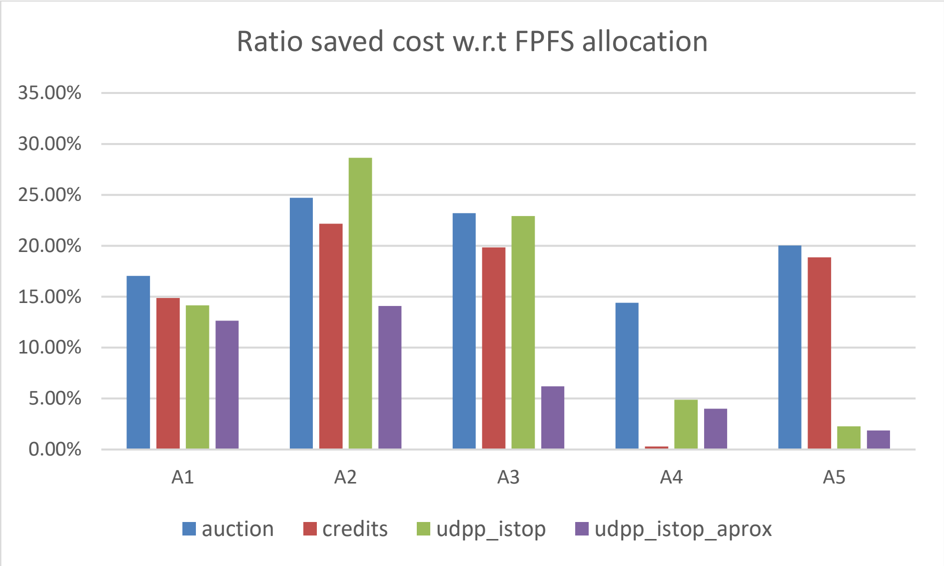

Figure 7. Ratio of saved costs w.r.t FPFS allocation disaggregated by airline

According to this metric, we should observe an equal improvement percentage for all airlines to consider that a mechanism meets the optimal level of equity. However, we observe the same trend as with the previous metrics.

The auction mechanism improves the situation of all airlines in a more even way. They all improve between 15% and 25%, which is not an excessively large difference. The credit mechanism follows the auction performance very closely, with the exception of airline A4, which practically does not improve its situation, as explained before. Finally, the udpp_istop and udpp_istop mechanisms present a greater variability of improvements between airlines. The reason is the same as the one detailed above.

Finally, it is important to stress that extremely strong conclusions should not be drawn from this equity analysis. The results could be sensitive to the rigidity of behaviours imposed on airlines along with the specific traffic and network used in the simulation. The work carried out in WP5 with a model much more similar to reality and different traffic conditions may serve to contrast these first intuitions.

5.4 Robustness

One of the ambitions of BEACON is to be able to measure how the performance of the flight prioritisation mechanisms varies when airlines behave in a "non-rational" or strategic manner, which is the case most of the time. Here, the **robustness is understood as the ability of a flight prioritization mechanism to maintain the same performance regardless of the degree of irrationality in the behaviour of the decision-making agents**.

The robustness of each mechanism has been assessed by comparing the results of some selected metrics in a rational scenario and the results of the same metrics for other scenarios with the chosen behavioural biases included (see sections 3.3.5 and 4.4). The results of these scenarios are shown in Figure 8 and Figure 9 along with the results for the 'rational' scenario for some of the most relevant metrics from the above KPAs.

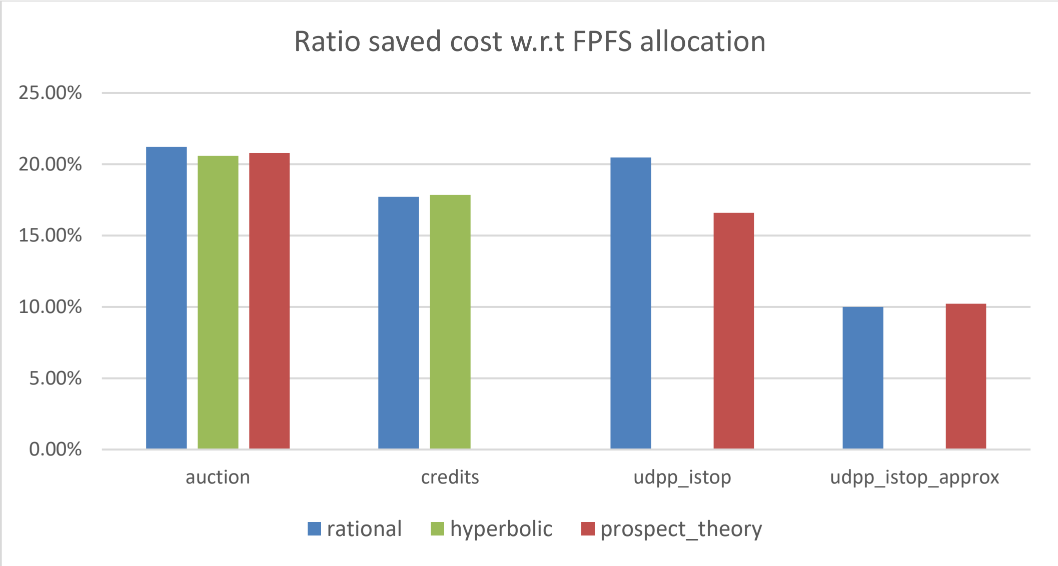

Figure 8. Ratio saved cost w.r.t FPFS allocation taking into account the behavioural models

Figure 8 again shows the percentage of cost saved by each airline with respect to the allocation based on FPFS but this time including the scenarios with behavioural biases. In general, the mechanisms seem quite robust against the irrational practices simulated in the model. At first glance, there is no clear trend, some mechanisms slightly worsen their performance (auction and udpp_istop) while others are not affected by the modelled biases or even minimally benefited (credits and udpp_istop_approx).

The auction presents a fairly robust behaviour. Although its performance worsens, the ratios of saved costs are minimally different from that calculated with the rational scenario, especially with prospect theory. It seems that these biases barely change the final allocation results, thus providing a very similar performance. In future work, it would be very interesting to test how this conclusion may vary by including more extreme and distinct airline behaviours.

On the other hand, the mechanism that presents a greater variation to the introduced behaviour bias, and therefore less robustness, is the udpp_istop. In this case, it seems that the fact that the airlines are not calculating the cost of the delay in an absolute way but based on the established reference point (see section 3.3.5.1.1) causes that many of the beneficial swaps between airlines do not occur. Consequently, the improvement provided by the mechanism is lower.

The situation of the credits and udpp_istop_approx mechanisms is a little bit different. Although both mechanisms present minimal changes in performance against 'irrational' practices, and therefore seem quite robust, this minimal change is favourable to the irrational scenario, meaning that, counterintuitively, the mechanisms offer slightly better results (≈ 0.2%) for scenarios where biases are active.

The main reason for this phenomenon lies in the approximation of the cost function. Both mechanisms are based on prioritisation process which relies on the reconstruction of the flight cost function by the Network Manager. If this function is not sufficiently adjusted to the real one, the mechanism reduces its performance, as described in the previous sections. For both mechanisms it seems that the variability produced by the different bias applied overcomes the poor approximation and some prioritisation opportunities that were missed with the rational scenario are computed now. However, this happens very rarely, so the improvement is minimal.

Finally, we can also appreciate the effect of these behavioural models by observing the value change for each airline. Figure 9 shows the same metric as above but disaggregated by airlines and for each mechanism. In general, no clear trend can be seen. For all the mechanisms we see that there are airlines that slightly improve their situation while others worsen it a little.

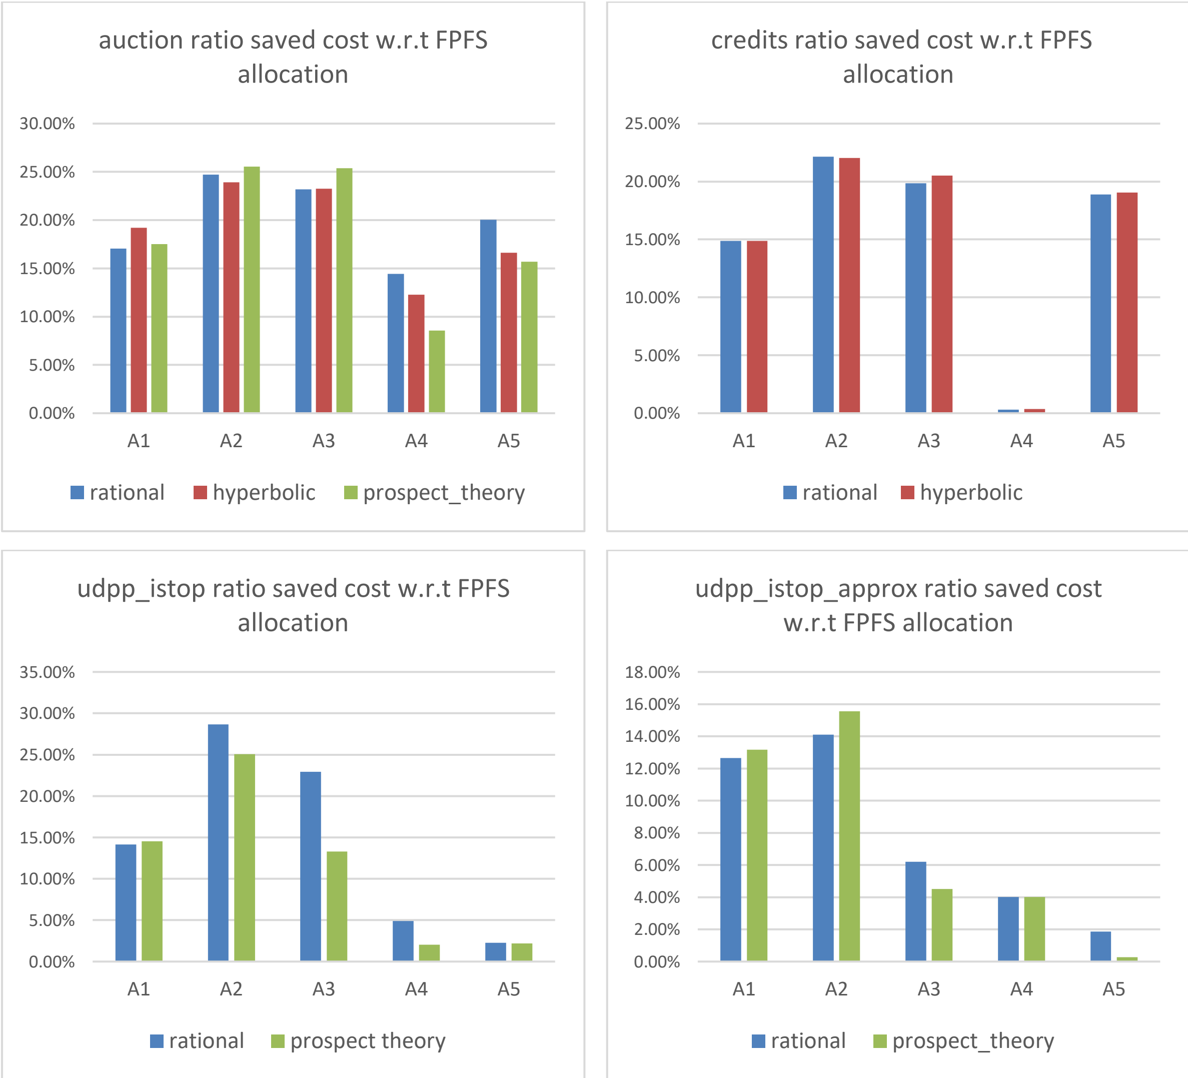

Figure 9. Ratio saved cost w.r.t FPFS taking into account the behavioural models disaggregated by airline

Again, it is important to note that this is a first analysis of the robustness of the mechanisms where it is possible that some results are influenced by the simulation hypotheses and the chosen scenarios. This should be completed by the work being done on WP5.

6 Conclusions and next steps

From the detailed analysis of the experiment results, some conclusions and insights can be drawn:

- The true potential of flight prioritization mechanisms to reduce airline costs is appreciated when regulations are large enough to involve delays potentially close to breaking operational margins for flights.
- All the mechanisms tested provide an improvement in terms of saved cost. This fulfils the main objective of a flight prioritisation mechanism: to provide airlines with methods with which to reduce delay costs associated with their flights affected by ATFM regulations. Among them, the auction mechanism stands out.
- The performance of the mechanisms where the Network Manager is in charge of optimizing the final sequence of flights based on certain operational parameters is highly dependent on the quality of the approximation of the cost function. This is clearly evidenced in the case of the udpp_istop_approx mechanism. Good approximations of the cost functions are required.
- None of the mechanisms reaches an optimal level of equity in the proposed terms. However, some mechanisms like the auction or credits offer quite good results. A sensitivity study to the model and the scenario conditions is necessary to confirm these results.
- In general, all the mechanisms offer a good level of robustness against the behavioural biases implemented in the model. The 'irrational' practices slightly worsen the performance of the auction and the udpp_istop mechanism. However, the credits mechanism and the udpp_istop_approx mechanism are minimally improved due to the poor approximation to the cost function.

The results shown are sometimes conditioned by the characteristics and the assumptions of the model and require further investigation, which could be addressed in WP5, if it fits into the scope, or in another future research project. The following questions remain to be explored:

- To which extent the current setup influences the results (to be explored in WP5)? In particular, will the trends on equity and cost-efficiency hold with more diverse regulation (number of flights, etc)
- What is the influence of the mechanism parameters, in particular the initial credit endowment and the prices of margin and jumps? (to be explored to some extend in WP5).
- What is the influence of the behavioural parameters (risk aversion, etc)? (to be explored most likely in a future project)

7 References and Acronyms

Applicable documents:

[1] Grant Agreement No 893100 BEACON – Annex 1 Description of the Action.
[2] BEACON Consortium Agreement
[3] BEACON D1.1 Project Management Plan, Issue 1, July 2020

Reference documents:

[4] Arthur, W. B. (2006). Out-of-equilibrium economics and agent-based modeling. Handbook of computational economics, 2, 1551-1564.
[5] Cardoso, C., Bert, F. & Podestá, G. Modelos Basados en Agentes (MBA): definición, alcances y limitaciones.
[6] Bonabeau, E. (2002). Agent-based modeling: Methods and techniques for simulating human systems. Proceedings of the National Academy of Sciences, 99 (suppl. 3), 7280-7287.
[7] Mueller, M. and A. Pyka (2016) "Economic Behaviour and Agent-Based Modelling", in Frantz, R., S.H. Chen, K. Dopfer, F. Heukelom and S. Mousavi (Eds) Routledge Handbook of Behavioral Economics.
[8] EUROCONTROL Central Route Charge Office (2021) Customer Guide to Charges
[9] European Union (2004). Regulation (EC) No 261/2004 of the European Parliament and of the Council of 11 February 2004 establishing common rules on compensation and assistance to passengers in the event of denied boarding and of cancellation or long delay of flights, and repealing Regulation (EEC) No 295/91. Available at https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32004R0261&from=ES
[10] University of Westminster (2015). European airline delay cost reference values Updated and extended values Version 4.1. Available at https://www.eurocontrol.int/sites/default/files/publication/files/european-airline-delay-cost-reference-values-final-report-4-1.pdf
[11] BEACON D3.2 Industry briefing on updates to the European cost of delay, Issue 1, September 2021
[12] BEACON D3.1 High-level modelling requirements. Issue 1, December 2020
[13] BEACON D6.1 Intermediate Concept assessment report, Issue 1, September 2021
[14] BEACON D4.1 Behavioural Model Parameter Calibration, Issue 1, September 2021
[15] PJ19.04 D4.7 SESAR Performance Framework (2019) 01.00.01

Acronym	Meaning
ABM	Agent Based Model
AIBT	Actual In-Block Time
AOBT	Actual Take Off Time
ATC	Air Traffic Control
ATM	Air Traffic Management
ATFM	Air Traffic Flow and Capacity Management
ANSP	Air Navigation Service Providers
AU	Airspace User
ECAC	European Civil Aviation Conference
ETA	Estimated Time of Arrival
FPFS	First Planned First Served
FDR	Fleet Delay Reordering
GIS	Geographic Information System
HD	Hyperbolic Discounting
ICAO	International Civil Aviation Organisation
ISTOP	Inter Airline Slot Swap Offer Provider
KTN	Knowledge Transfer Network
KPA	Key Performance Area
KPI	Key Performance Indicator
MTOW	Maximum Take-Off Weight
NM	Network Manager
OD	Origin-Destination
PT	Prospect Theory
SESAR	Single European Sky ATM Research
SES	Single European Sky
SFP	Selective Flight Protection
SOBT	Schedule Off-Block Time
UDPP	User Driven Prioritisation Process
WP	Work Package

Appendix A Flexible credits analysis over time

The operation of the credit-based mechanism is deeply influenced by the number of credits that the airlines have thought the simulation. When an airline has one or more flights affected in a regulation, it can use its credits to update the margin values and the jump of its affected flights. If this update requires the payment of credits, as detailed in section 3.3.4.2.1, and the airline does not have enough credits, the values of the parameters sent to the Network Manager will be far from reality and therefore the approximation of the cost function will be worse. As a consequence, the final optimization of the sequence based on the reconstructions of these cost functions will be suboptimal.

As a complementary part to the analysis of results, some specific details of the credit mechanism have been studied so we can draw better conclusions from the results. First, Figure 10 shows the credits available to each airline at the end of each simulation day.

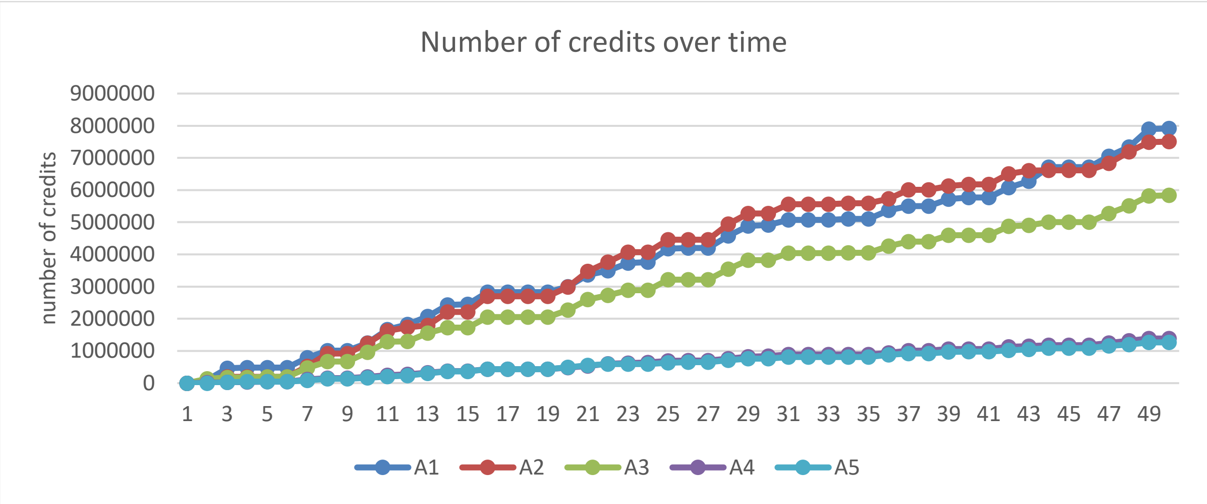

Figure 10. Number of credits over time per airline

The trend is clear, the number of credits from all airlines is increasing over time. This is mainly explained by the fact that, for these experiments, the default margin and jump values (see section 3.3.4.2.1) have been defined in such a way that airlines almost always earn credits by adjusting the parameters of their affected flights. It turns out that, in most cases, the default value chosen for the margin (60 min) is lower to the real first margin of the affected flight and the default value defined for the jump (50000 EUR) is higher than the real jump. Consequently, when the airline set the parameters to their real values it should increase the margin and lower the jump, receiving credits for it.

The number of credits that the airline receives throughout the day depends on the number of flights affected and their cost function, which is why there are differences between airlines, although they all increase the number of credits. This basically implies that the airlines always have enough credits to properly execute their prioritisation decisions.

However, this trend is not desirable from the operational point of view of the mechanism. In the real world, if an airline realise that is constantly earning credits and always have plenty to spare, it could start using those credits to alter the margin and jump parameters above their honest values. Although this could, a priori, benefit that specific airline locally, it would also have a very negative impact on the others, drastically reducing the cost efficiency and fairness of the mechanism. We cannot observe this behaviour in the results due to the way the airlines' decision process is designed in this mechanism (3.3.4.2.1): the airlines always try to change the margin and jump to the real values no matter how many credits they have.

WP5 will perform a better calibration of the margin and jump standard values seeking to stabilise the credit gain over time. In addition, some complexity will be added to the behaviour of the airlines to make it more realistic.

Finally, the performance of the mechanism over time has also been studied with the intention of seeing if the amount of airline credits influences performance. Figure 11 shows the ratio of the saved cost w.r.t FPFS allocation for each simulation day in which there is one or more than one hotspot. The mechanism performance shows ups and downs without any clear trend. There is no clear relationship between credits and performance. As detailed before, the airlines always have enough credits to make the decisions they want, so we can assume that under these circumstances the performance of the mechanism is only influenced by the characteristics of the traffic. It will be interesting to see in WP5 how this performance changes for a situation where airlines don't have enough credits.

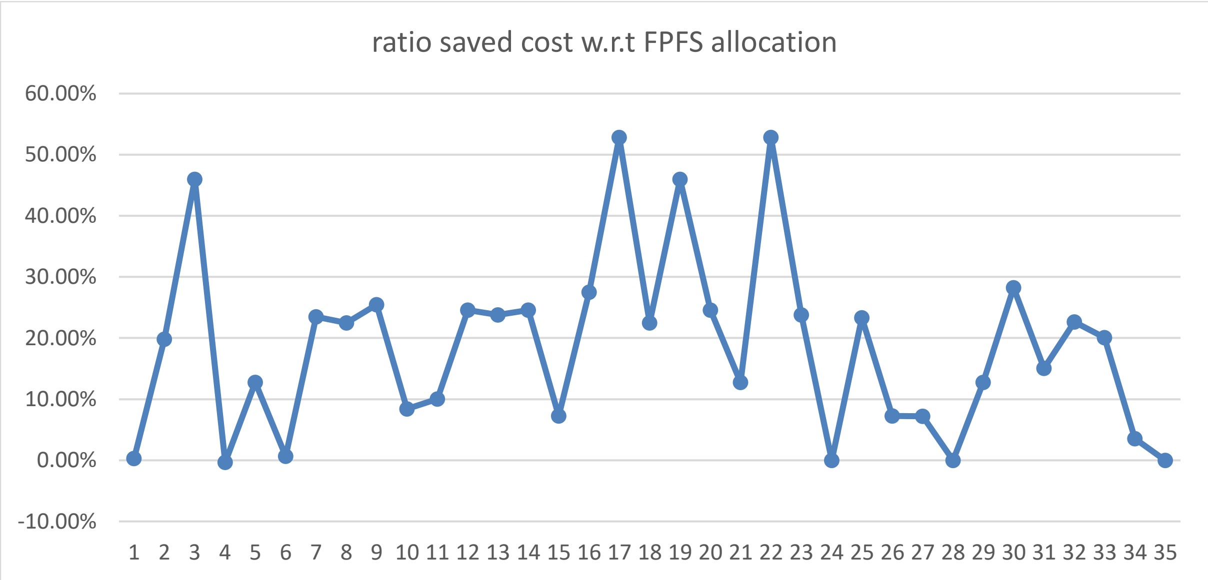

Figure 11. Ratio saved cost w.r.t FPFS allocation with the credit mechanism for each day.