


Study on the Modelling of Airport Economic Value

Final Report



Background information

This report presents the final status of the study on modelling Airport Economic Value.

This study was commissioned by the Airport Research Unit of EUROCONTROL to the University of Westminster supported by Innaxis as part of its contribution to SESAR Operational Focus Area (OFA) 05.01.01 entitled 'Airport Operations Management', in relation to the development of the AirPort Operations Centre (APOC) concept.

The objective of the study is to extrapolate the model used for capacity planning of en-route ACC centres to airports with a view to evaluate the optimum airport economic value, based on several cost functions and quality of service evaluations.

For any questions on the content of this report, please contact us at:

denis.huet@eurocontrol.int
cookaj@westminster.ac.uk

Authorship

Gérald Gurtner, Andrew Cook, Anne Graham – University of Westminster (London)
Samuel Cristóbal – Innaxis (Madrid)

Contents

1	Introduction and objectives	11
2	Literature and data availability reviews	13
2.1	Literature review	13
2.1.1	The implications of airport congestion and delays	13
2.1.2	Soft management options	14
2.1.3	Hard infrastructure approaches	15
2.2	Data availability	17
3	Model preparation and calibration	20
3.1	Preparing for the modelling process	20
3.1.1	Using the literature review	20
3.1.2	Using the data analysis	21
3.2	High-level principles of the modelling process	29
3.3	Calibration	31
3.3.1	Functional relationships	32
3.3.2	Direct calibration of parameters	33
3.3.3	Post-calibration of parameters	34
4	Model implementation and results	36
4.1	Implementation	36
4.1.1	Engine implementation	36
4.1.2	Graphical user interface (GUI)	36
4.2	Calibrated model results	38
4.2.1	Calibrated airport	38
4.2.2	Comparing different airports	40
4.2.3	Effect of parameters on the position of the optimum	42
4.3	Exploratory results	46
4.3.1	‘Better shopping time’	47
4.3.2	‘Longer shopping time’	47
4.3.3	Both effects	48
5	Conclusions	50

A	Data analysis	54
A.1	Correlation structure	54
A.2	Clustering analysis	61
A.2.1	Methodology	61
A.2.2	Analysis of clusters	62
B	The model	65
B.1	Calibration	65
B.1.1	The problem of delay	65
B.1.2	Delay and capacity	65
B.1.3	Cost of delay	67
B.2	Model implementation	69
B.2.1	Probability of operation	69
B.2.2	Implicit equation of delay	71
B.2.3	The airport	71
B.3	Exploratory analysis	72
C	Sensitivity analysis	74

List of Figures

1	Maximum of airport net income as a function of its capacity (shown in blue), and the corresponding average delay per flight (shown in red). These results arise from a calibration of the model on a major European hub.	9
2	Screenshot of the Graphical User Interface to the model.	10
3.1	Weights of the initial variables for each component.	24
3.2	Illustrative plot of the choice function. When the costs decrease in absolute value, the probability of operating the flight increases towards 1. In harsh decision conditions ($s \ll 1$) the airline stops operating the route as soon as the cost is non-zero.	30
4.1	Screenshot of the visualisation layer.	37
4.2	Help menu.	38
4.3	Evolution of the net income of the airport as a function of the capacity and the marginal operational cost.	39
4.4	All results are expressed per day. Evolution of (a) revenues per passenger, (b) cost per passenger, (c) net income per passenger, (d) average delay per flight, (e) cost of delay per flight, (g) ratio between operated flights and demand, (h) number of flights actually operated, (i) total net income of the airport, and (j) passenger utility, as functions of the airport capacity.	40
4.5	Threshold of the marginal cost of capacity for which an increase 50% of the capacity is profitable for different airports.	42
4.6	Optimal capacity as a function of the average load factor.	43
4.7	Evolution of different metrics with the increase of the standard deviation of delay at the airport (decrease of predictability).	44
4.8	Optimal capacity as a function of the standard deviation of delay. . . .	45
4.9	Different outputs of the models with satisfaction effect (see caption of Figure 4.4 for the meaning of each graph). These plots have been obtained with the following values of free parameters: $t_e = 0$, $s_e = 1500$, $cap = 24$, $\alpha = 40000$	47

4.10	Different outputs of the models with both effects (see caption of Figure 4.4 for the meaning of each graph). These plots have been obtained with the following values of free parameters: $t_e = 20$, $s_e = 1500$, $cap = 24$, $\alpha = 40000$	49
5.1	Illustration of optimum capacity at an airport.	52
5.2	Complex landscape with local minima arising from the non-trivial non-aeronautical revenues per passenger.	53
A.1	Number of communities in the network of airports as a function of the scaling parameter. The presence of plateaus indicates the existence of natural cluster structure at these scales.	63
A.2	Average Normalised Mutual Information (NMI) between partitions based on disturbed data and the original 3-clusters partition, as a function of the level of noise applied to the data. The error bars represent the standard deviations over the 100 realisations of noise.	64
B.1	Distribution of delays in CODA data ('real' delays) and DDR data for LFPG. The data has been truncated in the $[-15, 120]$ region (following truncation from CODA data).	66
B.2	Delay in seconds as a function of the traffic, for each hour of the day. The solid line is an exponential fit.	66
B.3	Histogram of the delays for LFPG at 9am, with a lognormal fit.	69
B.4	Correction of cost of delay for airline. The blue points are the cost of the average delay, the violet ones are the expected costs taking into account the distribution of delay, and the solid red line is a quadratic fit of the latter.	70
B.5	Evolution of the function w with the delay, with $\delta t_{init} = 9.5$, $w_{init} = 17.8$, $t_e = 15$, and $s_e = 1000$	73
C.1	Evolution of the average delay (left) and revenues of airlines (right) in the calibrated model for various values of the smoothness parameter s	74

Executive Summary

The primary objective of the Airport Economic Value project is to assess the value of additional passengers or additional capacity at an airport. It aims to qualify and quantify the main relationships and trade-offs between capacity, quality of service and profitability. This study provides a better understanding of the interdependencies of various KPIs and assesses the existence and behaviour of an airport economic optimum, in a similar way to the early 2000s, when estimating the economic en-route capacity optimum.

In order to do this, the project builds a functional model based on supply and demand curves. The implementation follows a data-driven approach. The modelling decisions are supported by a literature review and data analysis only; the latter encompasses multiple techniques from knowledge discovery, clustering and factor analysis, among others. Most of the more technical details have been presented in annexes.

This report presents the final model, the data analysis performed to support it, and a review of the literature and data availability. The aim is to present the work that led to the implementation of the model, the results obtained during the design process, the model itself, including its assumptions, and some key results obtained as model outputs.

The baseline year for the analyses is 2014. Operational and traffic data (from FlightGlobal and EUROCONTROL) and passenger data (from ACI EUROPE) both relate to this reference year. Obtaining 2014 financial data proved less straightforward than anticipated. FlightGlobal was still citing 2013 data for many airports. It was also necessary to extend the *depth* of the data, such that it was decided to use 2013 financial data from the Air Transport Research Society for all the analyses, as a proxy for 2014: ATRS only had 2013 data available at the time of the data analysis. For this initial model, it was only necessary to establish fundamental relationships between the data fields. Any obvious shortcomings regarding the financial data have been monitored and flagged.

‘Soft management’ and ‘hard infrastructures’ are considered for capacity increases. The literature review identifies many of the key potential trade-offs relevant to this research. It discusses the implications of delays at airports and the various approaches

different airports can adopt to increase their capacity. It considers key airport management issues such as non-aeronautical revenue generation and service quality.

The review also covers the complex area of airport charging/economic regulation and the wider consequences for airline network planning, identifying the difficulties involved with including such factors in any trade-off analysis.

Data availability is discussed. A small number of metrics and mechanisms are selected. Several data analysis techniques are identified which have been used to guide the modelling process, including principal components analysis (PCA) combined with cluster analysis. This unsupervised cluster analysis has confirmed, quantitatively, the usual qualitative distinction between airports, with primary and secondary hubs, etc. These clusters are then used in the model. The assumptions and the mechanisms of the model are described in detail, as well as the calibration process. To our knowledge, this is the first time that such a very wide range of data – in particular, economic and financial data – have been synthesised in one database and used to characterise airport performance.

The considerations from the literature review and data analysis lead to several conclusions. First, it is important to have a model which can be differentiated based on the type of airport considered. Second, the impact of delay on the airport results from airlines being less willing to operate a route because of the corresponding costs they incur. Third, uncertainty means that even levels of traffic under the theoretical airport capacity can imply some delay, and that the mean average delay is not a sufficient measure for consideration. Fourth, the modelling of the decision-making process of the airport cannot take into account changes in airport charges, but merely comparisons of the marginal operational cost of some extra capacity with the increase in demand due to decreased delays. Finally, the passenger perspective should not be directly modelled, because passenger choice of airport is largely independent of factors that can be influenced by the airport.

It is demonstrated that the model can be easily calibrated on real data and runs very well for the airports in the dataset. It produces reliable and realistic results. The fully calibrated results show the presence of a trade-off between the cost of extra capacity and the increase in the number of flights operated. As a consequence, airports usually have a maximum in their net income as a function of capacity, as shown in Figure 1. This maximum usually implies that the average delay at an airport is non-zero, i.e. that an airport operates slightly above its capacity. This is analogous to the situation for en-route delays. This airport delay can be further reduced by higher load factors and a better peak/off-peak traffic balance.

All the airports exhibit a maximum in net income as a function of capacity, if the marginal cost of operating extra capacity is sufficiently low. This threshold in the marginal cost is, however, rather different across airports, and only a few airports can




Figure 1: Maximum of airport net income as a function of its capacity (shown in blue), and the corresponding average delay per flight (shown in red). These results arise from a calibration of the model on a major European hub.

sustain a high cost of capacity: these are the largest and most congested airports, which clearly need extra capacity. This threshold is roughly consistent with the airports' current operational cost of capacity, which means that they should be able to manage this growth, subject to the availability of investment.

More exploratory results show that the picture can be significantly modified by the introduction of variable, non-aeronautical revenues per passenger. When tendencies to 'shop more with more time' and 'shop better with increased satisfaction' are introduced, the net income can exhibit different maxima and minima. The direct consequence is that an airport would probably not be willing (or able) to invest sufficient capital to reach the global maximum, and is likely to be 'trapped' in a local maximum. Since an increase in capacity is incremental (e.g. new runway, new terminal), this may actually render it impractical for the airport to reach any maximum.

The tool developed in this project could be used for various applications. First, with more detailed airline and passenger data, it could be used as a decision-making support tool for new investments at an airport. For example, the tool could easily show if new infrastructures are needed at congested airports, or if incentives to airlines should be developed for higher load factors. Moreover, it is easy to use and features a Graphical User Interface (GUI), as shown on Figure 2, which allows to explore interactively the model.

The tool could also be used by regulators to drive the adoption of new regulations for different groups of airports, based on the results of the model and the cluster




Figure 2: Screenshot of the Graphical User Interface to the model.

analysis.

It is also planned that the tool will be expanded and refined. In particular, the heterogeneity of business models among airlines needs to be taken into account in order to have a more reliable estimation of their reaction towards the cost of delay. Moreover, airlines base such decisions on at least two airports (destination and arrival), if not on several at once (assessing network effects and benefits). This needs to be taken into account in a broader model featuring different kinds of airports and their relationships.

Further collaboration with specific airports and ACI EUROPE are also planned. In particular, better and more detailed data will lead to tailored airport models, improving the prediction capacities of the present model.

1 Introduction and objectives

The primary objective of the Airport Economic Value project is to assess the value of additional passengers or additional capacity at an airport. It aims to qualify and quantify the main relationships and trade-offs between capacity, quality of service and profitability. The purpose of the project is to produce a model capable of capturing the consequences of the decisions taken by different types of airports, for example regarding possible expansion.

In order to do this, the project builds a functional model based on supply and demand curves and utility theory embedding causal relationships. Although the Airport Economic Value model started from a theoretical perspective, the implementation follows a data-driven approach. All modelling decisions are ultimately supported by data and data analysis only; this encompasses multiple techniques from knowledge discovery, clustering and factor analysis, among others.

Decisions at the airport are typically based on capacity considerations (both for airlines and passengers), management processes, profitability of added passengers or airlines, and future demand. These decisions result in multiple consequences for airlines and for passengers, which adapt their behaviour to the new environment provided by the airport. Moreover, the airport usually has to deal with considerable uncertainty regarding the demand, due to the long-term time frame of the decisions. Airport expectations are thus highly important for the model as they shape the future in which the consequences of their actions will be manifested. The final model of the Airport Economic Value project is able to answer very specific questions related to these mechanisms, such as the economic viability of the construction of a new runway at a given airport based on several parameters, such as the size of the airport and the requirements of the airport and airlines.

This deliverable presents the final model, the data analysis performed to support its creation, and the review of the literature and data availability. The aim is to present work which leads to the implementation of the model, results obtained during the design process, the model itself with all its assumptions, and some key results obtained from the model. The main text is intentionally concise. Further material can be found in the annexes, such as full tables and equations.

The report is organised as follows. Section 2 provides the literature and data

review. Section 3 then details the building process of the model, including the main results of the final data analysis. This is followed by Section 4 which presents a short overview of the technical implementation of the model, followed by some key results obtained with the model. We finally summarise the lessons learnt from the model, and its possible applications, in Section 5. The annexes contain the full technical details related to data analysis and the model.

2 Literature and data availability reviews

2.1 Literature review

Summary

This section reviews the academic and industry literature related to the Airport Economic Value concept by discussing the implications of airport delays and the options for capacity expansion. The main mechanisms of relevance are ‘soft management’ processes (such as price schemes, slot allocation strategies, etc.) and ‘hard management’ processes (such as infrastructure improvement). The main variables to consider include: capacity utilisation; traffic mix; aircraft occupation; ‘Net Basic Utility’ (revenues taking the cost of delay into account); the value of time of passengers; the share of domestic flights; and, all metrics related to the infrastructures of airports (e.g. the number of runways).

2.1.1 The implications of airport congestion and delays

Excess capacity will create minimal delays but will be unprofitable for airports, which will be incentivised to utilise their facilities as much as they can, since a significant proportion of their operating costs are fixed, giving a relatively low cost elasticity (0.27 (US) [1] and 0.3-0.5 (UK) [2]). Excess demand will produce delay costs for airlines and passengers [3]. Some airports will pay service-level rebates to airlines when delays and congestion occurs (4-7% of airport charges at Heathrow and Paris [4], [5]).

The delays will mean that passengers spend longer at the airport and, assuming that this translates into additional dwell time to use for commercial purchases, this can be viewed as a positive externality of congestion [6] This is supported by evidence that shows that there is a favourable influence of dwell time on passenger spending

[7], [8], [9] with time pressures having a detrimental impact [10]. However, a negative relationship between unit commercial revenues and passengers has also been found, arguably due to congestion discouraging sales [11] [12] [13]. No direct significant relationship between commercial revenues and delayed flights has been found [11] although it has been argued that congestion in the terminal should be differentiated from congestion on the runways, since passengers cannot control their time to take-off or shop whilst the aircraft is queuing, whilst they can choose when to arrive at the airport [14].

Airport passenger satisfaction, which is likely to drop with congestion and delays, may also influence commercial spend. At a global level a 1% increase in passenger satisfaction scores is associated with a 1.5% growth in commercial revenues [15]. So even though greater satisfaction may not directly influence passenger airport choice (this being driven more by locational factors and airline fare/service preferences [16]) it may help airport profitability by enhancing commercial revenues. The resulting relationship between satisfaction and profitability has not always been confirmed [17] as research here is very scarce because of the lack of appropriate and publicly available satisfaction data.

2.1.2 Soft management options

In trying to match more closely demand and capacity, the literature discusses two main options for airports. First, there are so-called ‘soft’ management approaches, that tend to be quick to implement, potentially low cost, but limited in scope as they do not involve any major changes to the physical infrastructure. ‘Hard’ options, by contrast, are slow to implement and expensive. These can yield large increases in capacity, because they are lumpy and are made infrequently in relatively large, indivisible units. These two approaches are simultaneously considered by airports, and lead to a two stage optimisation, as shown in [18], with different time frames.

The soft options can relate to both strategic planning and tactical adjustments [19]. In the broadest sense, these can include substituting short-distance air travel with high-speed trains, diverting traffic to other airports or using multi-airport systems [20]. Related to the airport itself, options may be infrastructure improvement planning [21], [22], changing the ATC rules, and reorganising traffic to make better off-peak use of facilities, or by using aircraft with higher seat capacity, even though this may lead to additional congestion in the terminals [23], [24].

On the demand side, a major consideration is whether congestion or peak pricing can be used to manage the traffic. This has been discussed in depth in theory [25] particularly related to the ability of dominant carriers with market power (rather than being atomistic in perfect competition) to internalise the congestion costs. Potentially, peak pricing would push up airline and passenger costs, and bring extra revenues to the airport, (assuming that it is not introduced in a revenue-neutral manner).

However, in practice, peak pricing has proved difficult to implement and is unpopular with airlines, particularly since it is viewed as unfairly discriminatory and considered ineffectual in changing behaviour because of complex scheduling operations and slot allocation constraints [26].

Relating congestion pricing to passenger value of time, theoretical research has demonstrated that business passengers, exhibiting a greater value, would benefit from higher charges to protect them from excessive congestion caused by leisure passengers with a lower relative value of time [27] [14] and the optimal airport charge would be higher than if passengers were treated as a single type [6]. However in practice airports do not discriminate between business and leisure passengers in their pricing, and even using airline models types as a proxy (e.g. full service carrier - FSCs, low cost carriers - LCCs) does not hold really true in today's environment as the distinction between the markets for these two types of airlines is becoming increasingly blurred.

In the short-term, any changes in prices to reflect congestion may not be possible if the airport is subject to economic regulation, especially incentive regulation, which typically places limits on the price increases which are allowed [28]. An alternative demand management technique, frequently researched [29] [30] and independent of the economic regulation mechanism, is a reformed slot allocation process, probably using slot auctions or trading systems, which would have major financial consequences for airlines and passengers, but less certain impacts on airport revenues.

2.1.3 Hard infrastructure approaches

In discussing the provision of hard infrastructure, it has been argued that the uncertainty of future demand [18] and the unpredictability of capacity degradation should be considered [31]. Increasing local capacity can have major unforeseen wider impacts, for example, because of the network effects of delays [32]. Trade-offs between providing different types of capacity at departure, and capacity at arrival, have been identified [33] and the relationship between runway and terminal capacity examined [14], [34]. It has been contended that the runway capacity should be prioritised since this is what causes bottlenecks for most airports [23], [24]. There is also the trade-off between focusing on operational and commercial capacity and the extent of complementarity between these two different areas. This is affected by the choice of till or cost allocation method used when setting prices, with the single till including all airport activities compared to the dual till, when just the aeronautical aspects of the operation are taken into account [35], [6].

Airports will have different incentives to invest, particularly if they are subject to economic regulation. So-called 'cost-based' or 'rate-of-return' regulation can set incentives for excessive and too costly investment; price-cap regulation, whilst providing incentives for cost efficiency, can be associated with under-investment. The situation with light-handed regulation is not so clear [36].

If the airport does invest, growing in size and evolving into a new type of airport with different operations and/or traffic mix, there will be cost and revenue implications. Larger airports are generally able to provide a greater range of commercial facilities and services, increasing the commercial spend (less than US \$5 per square metre for airports of less than 5 million passengers growing to excess of US \$30 for airports with more than 25 million passengers [37]. Leisure passengers have been shown to spend more than business passengers [11], [9] with some evidence indicating that LCC passengers spend less [9], [38].

Traffic mix changes will also bring associated costs, related to the service expectations of the airlines, such as ensuring a fast transfer time for hub airports, or swift turnarounds for LCCs. As regards airport size, evidence is mixed but generally it shows that airports experience economies of scale, albeit with different findings related to if, and when, these are exhausted and if diseconomies then occur [39], [40], [41]. As a consequence of these apparent cost and revenue disadvantages for smaller airports, the European Commission (EC)'s view is that airports under 1 million passengers find it hard to cover all of their operating costs, let alone their capital costs. At a size of 3-5 million they should be able to cover all their costs to large extent, whereas beyond 5 million they should be profitable [42].

The costs of any additional capacity can be allocated to airport charges in different ways. There may be a degree of pre-financing, unpopular with airlines, when certain costs are covered in advance of the capacity becoming operational [43]. Research also confirms the strong influence of market-oriented factors (price sensitivity, competition) on pricing [44], [45]), [46], [47]. The airline's responses to charge increases will depend on their relative importance to overall costs, with LCCs arguably being the most sensitive [48]. If charges are passed directly to the passengers, their own sensitivity will reflect the typically quite small influence of charges on airline fares, and subsequently their charges elasticity will be relatively small (estimated at Stansted to be less than -0.15, rising to -0.2 to -0.6 after considering some degree to airport substitution [49], suggesting a fairly marginal impact). However, the evidence is unclear as to the extent to which airlines pass on changes in charges, or whether they choose to absorb at least some of these, with a supply-side response to adjust capacity by making changes to routes and schedules [50]. Irrespective of whether charges are passed on or not, airlines may see their profit margins reduced, because any supply response will involve lumpy reductions in airline capacity, having considerable impacts on the passengers [51]. This is an example of the wider impacts of capacity provision, which, as with delay impacts, can intuitively be understood but are difficult to support with empirical research.

In summary, this literature review has identified many of the key potential trade-offs relevant to this research. It has discussed the implications of delays at airports and the various approaches different airports can adopt to increase their capacity. It has considered key airport management issues such as non-aeronautical revenue generation and service quality. This has helped to provide the research context, and to inform

the model where comparisons are made of the marginal costs of extra airport capacity with the increase in demand due to decreased delays. The review has also covered the complex area of airport charging/economic regulation and the wider consequences for airline network planning, identifying the difficulties involved with including such factors in any trade-off analysis. The overall discussion has enabled an assessment to be made of the main variables commonly used in the literature, for example related to aircraft movements, passengers, airport characteristics and capacities, which has informed our own choice of parameters. This now leads on to the consideration of data availability.

2.2 Data availability

Summary

Eleven sources of data have been considered and acquired. Data management included data cleaning and small extrapolations, as well as cross-checks for plausibility. Data acquisition and consolidation comprised a large effort within the project.

One of the most intensive efforts of the project has been dedicated to data acquisition, cleaning and consolidation. In order to have a model which could be calibrated as much as possible, different kinds of data have been collected. Table 2.1 shows a summary of the data collected, with a brief description of their use.

The reference year for the analyses is 2014, this being the most recent year for which the data required were most generally available. A major component was airport financial and operational data sourced (through subscription) from FlightGlobal (London, UK). ATRS (Air Transport Research Society; USA and Canada) benchmarking study data were purchased, in addition, particularly for the provision of complementary data on airports' costs and revenues. At the time of analyses, only ATRS data for 2013 were available, and these selected data were used as a proxy for 2014. Financial and operational data were compared with in-house, proprietary databases, with adjustments made as necessary. Data on airport ownership, and additional data on passenger numbers, were provided by Airports Council International (ACI) EUROPE (Brussels). European traffic data were sourced from EUROCONTROL's Demand Data Repository (DDR) with delay data primarily from the Central Office for Delays Analysis (EUROCONTROL, Brussels). Note that, importantly, *local* turnaround delay is used throughout this work, as this reflects airport *in situ* effects, whereas air traffic flow management departure delay is generated due to en-route delay, or delay at the destination airport - i.e. it is attributable to remote effects. We

Source	Typical Content	Use
FlightGlobal	Number of flights, number of passengers, share of European flights	Cluster analysis, calibration
EUROCONTROL CODA	Delay per airport & per type	Comparison with DDR delays
EUROCONTROL DDR	Full trajectories of aircraft for one month of data	Delay distribution, capacity fitting, share of different types of companies
ACI	Number of passengers (domestic, international, etc.)	Calibration purposes
ACI	Ownership airport	Not used in final analyses
Private communication, EUROCONTROL (2016)	Coordination of airport	Not used in final analyses
Skytrax, etc	Passenger satisfaction	Cluster analysis
ATRS	Financial data	Cluster analysis, calibration
ATRS	Airport charges	Comparison with aeronautical revenues per aircraft
Private communication, EUROCONTROL (2016)	Maximum Take-Off Weight	Cost of delay calibration
University of Westminster [3]	Cost of delay	Cost of delay calibration

Table 2.1: Data sources, content, and use.

did not have access to clean, local (airport generated) air navigation service (ANS) delay data. Other in-house sources of data were used in addition to those listed, also drawing on the literature review.

Considering the wider context of operations in 2014, there were 1.7% more flights per day in the EUROCONTROL statistical reference area, compared with 2013. The network delay situation remained stable compared to 2013, notwithstanding industrial action, a shifting jet stream and poor weather affecting various airports throughout the year, particularly during the winter months [52]. The average delay per delayed flight demonstrated a slight fall relative to 2013, and operational cancellations remained stable *ibid*. We return to the issue of industrial action shortly.

In the absence of access to a single, comprehensive source of passenger quality of service data, airports were assigned an overall passenger satisfaction ranking for 2014, initially based on Skytrax “The World’s Top 100 Airports in 2014” ranking data¹, and then adjusted according to independent reviews by two experts, in addition to some limited inputs from ACI (Montreal, Canada) drawing on its Airport Service Quality [15] programme data. On this basis, the airports were allocated to a ‘top’, ‘middle’ or ‘lower’ ranking. Notwithstanding fairly extensive industrial action in 2014², clearly impacting a number of passengers at specific airports, it is difficult to assess the collateral (confounding) impact of such events on corresponding passenger satisfaction scores for such airports. The final rankings derived cannot be shown due to confidentiality restrictions. This new parameter derived by the team is one of many important inputs informing the cluster analysis of 3.1.2.3.

To our knowledge, this is the first time that such a very wide range of data has been synthesised in one database and used to characterise airport performance.

¹http://www.worldairportawards.com/Awards/world_airport_rating_2014.html

²Air traffic control – Belgium: June, December; France: January, March, May, June; Greece: November; Italy: December. Airlines – Air France: September; Germanwings: April, August, October; Lufthansa: April, September, October, December; TAP Air Portugal: December.

3 Model preparation and calibration

The modelling process is presented in this section, as well as the resulting model and the calibration process.

3.1 Preparing for the modelling process

3.1.1 Using the literature review

The literature review has been used to inform the selection of the main mechanisms and variables for the final model. These are detailed below:

- Airlines are affected by delays through compensation payments and duty of care, as required by Regulation 261 [53], and through the loss of market share arising from reduced punctuality.
- Passenger spending is not directly dependent on the delay at the airport, but on overall passenger satisfaction.
- Airports create delays primarily by operating over or near their capacity thresholds.
- Many airport charges are subject to economic regulation and thus charges cannot generally be considered as a variable that the airport is freely able to adjust.
- Airport capacity is not a simple value, but rather a concept embedding several complex and interrelated mechanisms (terminal, runway, gates etc.) and uncertainty.
- Airports are quite diverse in terms of size, business models, types of airlines and passenger profiles.
- Passenger demand drivers are exogenous to the airport.

These considerations lead to several conclusions. First, it is important to have a model which can be differentiated based on the type of airport considered (hub airport, regional airport, etc.). Second, the impact of delay on the airport results from airlines being less willing to operate a route because of the corresponding costs they incur (and possibly as a shortfall in traffic too). Third, uncertainty means that even levels of traffic under the theoretical airport capacity can imply some delay, and that the mean average delay is not a sufficient measure for consideration. Fourth, the modelling of the decision-making process of the airport cannot take into account changes in airport charges, but merely comparisons of the marginal operational cost of some extra capacity with the increase in demand due to decreased delays. Finally, the passenger perspective should not be directly modelled, because passenger choice of airport is largely independent of factors which can be influenced by the airport.

3.1.2 Using the data analysis

Summary

The data analysis supporting the modelling process is presented. A variable correlation matrix is studied, giving some insights into the relationships between ‘extensive’ and ‘intensive’ variables, infrastructure and traffic data, financial and traffic data, etc. In order to reduce the complexity of the analysis, a principal components analysis shows that the main driver for variance is the size of the airport, followed by the airport business model and the balance between aeronautical and non-aeronautical revenues. A clustering analysis divides the main airports into three communities, which are roughly consistent with expert judgement.

3.1.2.1 Correlations

One of the challenges of constructing a comprehensive model for an airport is in building causal relationships between a small number of core variables. The choice of these core variables is determined by considering the dependencies of the different variables on each other in the data. The first step is thus to compute the correlation coefficients between each variables as these give the magnitude of the linear statistical correlations. The table of correlations can be found in Annex A.1, as well as the different variables considered and their meaning.

Airport operating revenues are strongly correlated with several metrics, including the number of passengers and the number of flights, which is expected, but also with aircraft occupation (number of passengers per flight) and the number of passengers per route, with correlation coefficients as high as 0.97. This is especially striking because the latter metrics are not trivially linked to the number of passengers or flights, so it is not a simple scaling effect. In fact, it shows how ‘extensive’ variables, i.e. scaling with the number of passengers or flights, can interact with ‘intensive’ variables. These effects are very important to capture, because intensive variables usually reflect the fundamental organisation of the system, related to the interaction between different agents (e.g. some kind of management rule). Regarding the precise meaning of this correlation, it is not clear at this stage why the operating revenues should be so closely related to these metrics, except if they are linked to some kind of capacity, as discussed below.

More interestingly, some of the intensive variables, like aircraft occupancy, are correlated with the size of the airport (0.61). This is also expected since small airports usually have more versatile functionality, which requires smaller aircraft for flexibility. Other features are worth exploring. For instance, the proportion of intra-European flights seems to be strongly (negatively) correlated with different variables, including the total number of flights and number of gates (-0.77 and -0.65, respectively). This was expected since intercontinental airports are also the largest ones. More importantly, total delays seem to be positively correlated with the number of runways, the number of gates and the number of terminals (0.41, 0.58, and 0.45, respectively), i.e. with the size of the infrastructure. This is because longer delays are expected at the bigger airports, which have the largest infrastructure. On the other hand, the delay per flight is less correlated with the infrastructure (0.38, 0.34, and 0.2). This is a good sign, because it could mean that the airports increase their infrastructure to counter-balance delays. It is also observed that runway and terminal usage have non-trivial behaviour with respect to the *number* of runways and terminals, since they are weakly or negatively correlated with them (-0.25 and 0.02), which could loosely mean that average airports are ‘over-building’, i.e. the number of runways and the number of terminals increase more quickly than the number of passengers. Note that, strangely, the terminal usage increases (weakly) with the number of runways (0.25), whereas the runway usage is quite independent of the terminal usage. This is the product of a subtle co-evolution of different capacities, namely the terminal capacity and the runway capacity.

Indeed, this is typically where simple correlation scores begin to show their limit. It is not clear at this point what are the drivers of the different metrics and whether a few causes only can explain most of the correlations. In order to explore this, we turn now to principal components analysis (PCA).

3.1.2.2 Principal components analysis

Since there are different types of airports, this is likely to have major consequences that should be reflected in the model. Rather than relying on purely expert-driven clusters, the clusters are defined in a data-driven way. In order to do this, a few variables of importance in the data were selected, which are presented in Table 3.1. We tried to select different types of variables with the most reliable data. Clearly, some of the variables are correlated, as shown in the correlation table in Annex A.1. In order to have a better picture, the decision was made to use PCA, reducing the number of independent variables to four, while keeping 80% of the initial variance.

Abbreviation	Short description
AO_tot	Number of airlines
CUI	Capacity utilisation index
NBU	Net basic utility (net operational income minus costs of delay)
cap	Runway hourly capacity
cht	Share of low-cost carriers and charters
fsc	Share of full-service carriers
delay_per_flight	Delay per flight
delay_tot	Cumulative turnaround delay
exp_tot	Total yearly expenses
flight_EU	Share of European flights
flight_per_rnwy	Flights per runway
flight_per_term	Flights per terminal
flight_tot	Total number of flights
gate_tot	Number of gates
term_tot	Number of terminals
pax_per_flight	Passengers per flight
pax_tot	Number of passengers
rev_areo	Aeronautical revenues
rev_non_area	Non-aeronautical revenues
rnwy_tot	Number of runways
route_tot	Number of routes
sat	Passenger satisfaction

Table 3.1: Variables used to characterise the airports.

Indeed, the objective of the PCA is to explain as much variance as possible in the data, and this is generally a key indication of the quality of the solution. However,

it is not acceptable to obtain a purely ‘mathematical’ solution in the analysis, i.e. whereby the analyst is not able to assign real meaning to the factors, which may be a challenge when there are too many of them. There is thus usually a trade-off between the number of components and the amount of variance explained.

It is also often desirable to ‘rotate’ the factors, to increase loadings on some of the original variables, and decrease them on others, in order to ease the interpretation of the solution and improve its simplicity. Thus to allow for a better interpretation of the results, we used varimax rotation. This is an orthogonal rotation method that minimises the number of variables with high loadings on each factor [54].

After this procedure, the four new components – linear combinations of the initial variables – have the weights displayed in Figure 3.1. The new variables explain

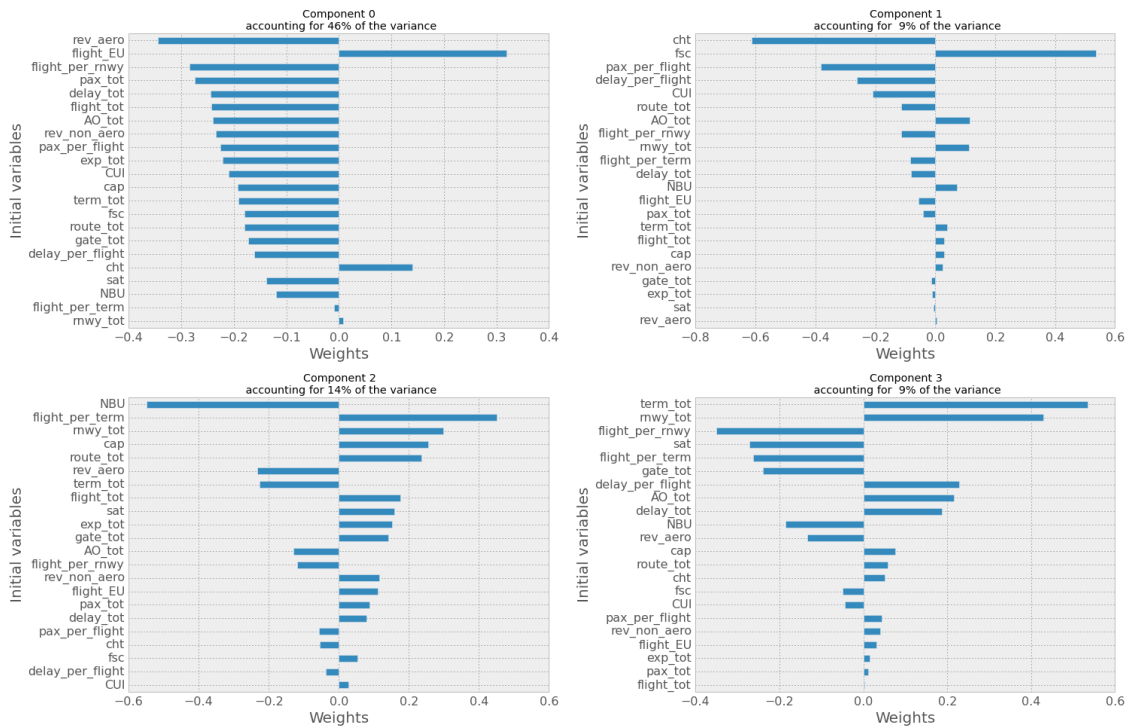


Figure 3.1: Weights of the initial variables for each component.

approximately 46%, 9%, 14% and 9% of the variance, respectively. The remaining variance (22%) can be explained by additional components, but the more we have, the less they are individually significant. For instance, the next component only explains 5% of the variance. Presenting four components, explaining some 80% of the variance of ~30 initial components, is sufficient to draw reasonable conclusions, even if their interpretation is not completely straightforward.

The first component (labelled 0) is homogeneously composed of all initial variables, in particular the ‘extensive’ variables such as the total number of flights or the number of delays. Hence, this first variable can be seen as the ‘size’ of the airport, which appears to be the main driver of most of the initial variables, because this first component accounts for almost half of the variance.

The second one (labelled 1) is clearly linked to the type of airlines which are operating at the airport. Specifically, it seems that 9% of the variance is closely related to the fact that airports serve more traditional/full-service carriers or more low-cost carriers. It is also evident that the infrastructure is closely linked to this, since the number of runways and terminals play a large role in this component too. Interestingly, the component is related to the number of passengers per aircraft, which is low when the component is low, i.e. when the airport is more ‘low-cost-oriented’ – in spite of pressures on these airlines to be punctual and have minimal turnaround times. This is not unexpected, since low-cost carriers often operate smaller aircraft, especially as they have very little long-haul traffic. It is also worth noting that the delay per flight increases when the airport is more ‘low-cost-oriented’.

The third component (labelled 2) is related to the financial state of the airport, with net basic utility and non-aeronautical revenues playing major roles. Interestingly, the total number of runways has a positive impact on this component, whereas the number of terminals has a negative effect. Since the capacity being measured is linked to the air traffic movements, it is clear that it has an impact on the component accordingly with the number of runways.

Finally, the last component (labelled 3) is linked to the physical infrastructure of the airport, which affects its usage (number of flights per runway and per terminal), but also the passenger satisfaction.

3.1.2.3 Clustering

Having defined a smaller number of variables with a clearer understanding of their meaning, we now cluster the airports, gathering together the ones which are similar in the same ‘community’. Note that in this part, the number of airports considered is equal to 32, which are the airports for which **all** the fields of table 3.1 are informed. Together, they represent more than half of the traffic in Europe in terms of passengers (53%).

There are many different ways of clustering data, depending on the definition of ‘clustering’. Several methods are routinely used in the literature but the specific choice of method is always quite subjective.

Indeed, the problem of clusterisation in community detection is mathematically poorly defined. The main reason is that the definition of a ‘cluster’ or ‘community’ is subjective, and varies among the different fields in which it is used. As a result, one should first define what one thinks would be the ‘right’ definition of a cluster

in the particular context, before trying to find the right method. In this research, we are interested in discovering whether some airports have similar behaviours. The simplest and most accurate way of finding this, it is proposed, is through the use of a Euclidean distance of their different characteristics, weighted by the PCA results. The best definition of a cluster is then related to the probability of two airports being closer (in terms of this Euclidean distance) to each other, than to the rest of the network. This technique, from network theory (and based on modularity), thus presents itself as being appropriate.

The details of the method can be found in Annex A.2. Here, we emphasise the fact that the previous PCA was directly integrated with the clustering analysis, since we used as distance between airports the Euclidean distance of the four components of the PCA, each weighted with their ratios of variance explained. In the same annex, we also show how we checked that the partition is robust with respect to uncertainty in the data. This partition is presented in Table 3.2. Cluster 1 includes mostly major hubs, whereas clusters 0 and 2 include airports with less traffic. Cluster 2 contains a number of secondary hub airports.

In order to inspect the clusters more closely, we also show in Table 3.3 the average value of each of the airports' characteristics, according to three categories: low, medium and high. Upon inspection of the table, the difference between clusters 0 and 2 appear more clearly. Indeed, the first one includes airports which have proportionally lower delays per flight, fewer routes, lower passenger satisfaction, fewer flights, and less congestion (lower CUI value) with respect to cluster 2. The table also confirms the status of 'major hubs' of the airports of cluster 1, with high numbers of passengers, numbers of flights, revenues and expenses. It confirms a tendency of such hubs to attract non-low-cost carriers, to experience higher delays per flight, and to have a more international profile. Interestingly, the passenger satisfaction is also different in this cluster, although the average level cannot be disclosed (for reasons of confidentiality). The net basic utility is not so high, however, probably driven by higher delays per flight, whereas the load factor is high for major hubs, as expected.

The cluster analysis thus gives us a suitable basis on which we can build differentiated models. In the following analysis, we use it to simplify some parts of the model. Taking into account the full diversity of the airports is not necessary.

Cluster Id	ICAO Code	Airport Name
2	EBBR	Brussels
	EDDL	Dusseldorf
	EGCC	Manchester
	EIDW	Dublin
	EKCH	Copenhagen Kastrup
	ENGM	Oslo Gardermoen
	ESSA	Stockholm Arlanda
	LOWW	Vienna
	LPPT	Lisbon
1	EDDF	Frankfurt
	EDDM	Munich
	EGKK	London Gatwick
	EGLL	London Heathrow
	EHAM	Amsterdam Schiphol
	LEBL	Barcelona-El Prat
	LEMD	Madrid Barajas
	LFPG	Paris Charles de Gaulle
	LIRF	Rome Fiumicino
	LSZH	Zurich
0	LTBA	Istanbul Ataturk
	EDDH	Hamburg
	EDDK	Cologne Bonn
	EFHK	Helsinki
	EGBB	Birmingham
	EGSS	London Stansted
	ELLX	Luxembourg
	EPWA	Warsaw Chopin
	LFMN	Nice Cote d'Azur
	LGAV	Athens
	LHBP	Budapest
	LKPR	Prague
	LPPR	Porto

Table 3.2: Composition of the partition from the clustering analysis. These airports combined represent 53% of the traffic in number of passengers.

	0	1	2
AO_tot	L	M	M
CUI	L	M	M
NBU	H	M	H
cap	L	H	M
cht	M	L	M
delay_per_flight	L	M	M
delay_tot	L	M	L
exp_tot	L	M	L
flight_EU	H	M	M
flight_per_rnwy	L	M	M
flight_per_term	L	M	L
flight_tot	L	H	M
fsc	M	H	M
gate_tot	L	M	L
pax_per_flight	M	H	M
pax_tot	L	M	L
rev_aero	L	L	L
rev_non_aero	L	M	L
rnwy_tot	L	M	L
route_tot	L	M	M
sat	*	*	*
term_tot	L	L	L

Table 3.3: Average value of each characteristic within each cluster (L=Low, M=Medium, H=High). The average level of satisfaction cannot be disclosed for confidentiality reasons, but are different in each cluster.

3.2 High-level principles of the modelling process

Summary

Important mechanisms and variables are considered and justified in this section. Airport congestion creates delay for the airline, and thus economic loss. An airline may decide not to operate from the airport if the cost is too high. The airport may increase capacity, which decreases the delay, but has an operational cost. The revenues of the airport are dependant on the number of flights (aeronautical revenues) and the number of passengers (non-aeronautical revenues).

In this section, the model used for the final version is presented. The corresponding equations, with some more details, can be found in Annex B.2. The model is a simple functional model based on representative agents, with the following fundamental principles.

First, we use a functional relationship linking the delay at an airport with the capacity and traffic. The progression is exponential (based on regression, see the calibration Section 3.3). The delay is then converted into a cost for the airline, based on the average maximum take-off weight (MTOW) at the airport. The cost is quadratic with the delay [3], which means that longer delays have proportionally higher costs than shorter ones.

This cost then fixes the probability that the airline actually operates the flight, through a probability function based on a hyperbolic tangent function. This choice is motivated by the fact that this probability is linked to some form of utility function for the airline, taking into account other (strategic) parameters (as described above). It allows us to have a smooth function which varies continuously between 0 and 1, and to have a ‘risk aversion’ of the agent which can directly be linked to the parameter s – henceforth referred to as the ‘smoothness’ of the decision. Indeed, when s is sufficiently small, the airline takes ‘harsher’ decisions, switching from operating to non-operating the route once costs are driven high enough. This behaviour is illustrated in Figure 3.2. Note that, in fact, we should strictly be referring to the net revenue contribution to the network, since airlines will tolerate loss-making legs that have a net benefit to the system. This parameter is the only behavioural parameter of the model and the only free parameter.

Note that the probability depends on the delay, which depends on the traffic, which in turn depends on the probability itself. As a consequence, there is a need to solve




Figure 3.2: Illustrative plot of the choice function. When the costs decrease in absolute value, the probability of operating the flight increases towards 1. In harsh decision conditions ($s \ll 1$) the airline stops operating the route as soon as the cost is non-zero.

an implicit equation in order to compute the delay (more details on this are to be found in [B.2.2](#)).

Once the exact number of flights has been found, together with the corresponding delay, the revenues of the airport are computed. There is an aeronautical part, proportional to the number of flights, and a non-aeronautical part, proportional to the number of passengers. The aeronautical revenue is simply the number of flights actually operated multiplied by the revenue per flight, computed from data, and is constant (although in reality aeronautical revenues are more complex being driven by the number of flights, weight of the aircraft and passenger numbers). The second one is the number of passengers – which in the model is the number of flights multiplied by the average load factor, extracted from data – multiplied by the revenue per passenger (also from data). This is fixed, too, as the (albeit limited) literature review showed that passengers do not seem to spend more at the airport if they are delayed, only if they have higher overall satisfaction, which is a constant here. Overall, it is important to note that **the total revenue per passenger is fixed** in the model. Only the number of flights change.

The net income of the airport is then obtained by subtracting the costs from the revenues. The cost has a fixed part, which is directly taken from data (total operational cost), and a part which changes with the capacity. For the latter, we use a simple linear equation with the marginal **operational** cost of extra capacity. Note

that this is **not** the cost of the infrastructure itself (e.g. runways, terminals), since we are interested in the recurring costs only, and not fixed lump sums associated with major investments.

Finally, in order to have an idea of the impact of delay on passengers, we compute a utility function based on value of time, which is a function of the share of business passengers at the airport. As an approximation for this, we use the share of traditional/full-cost carriers at the airport, as opposed to low-cost, although it is acknowledged that a growing share of business travellers now use low-cost carriers.

The model is not only based on average values, since one of the key features of the cost of delay is its non-linearity with the delay. As a consequence, we use certain distributions for the traffic and integrate the other values of the model with the traffic. More details about the types of distributions that we use are given in the following section on calibration.

3.3 Calibration

Summary

The model's calibration is presented. Data are used in three ways: (i) directly estimating input parameters for the model; (ii) extracting functional relationships from the literature or data regression; (iii) tuning input parameters by matching an output metric of the model to a target value from the data. The fully calibrated model has only two free parameters: the marginal operational cost of additional capacity and an internal behavioural parameter. The results are not strongly dependant on the latter, and it could in principle be post-calibrated using the average delay returned by the model. The first parameter should be considered as a variable, or be precisely estimated for individual airports.

Here we describe how we calibrate the model. The calibration itself is performed in various steps. Indeed, some parameters can be calibrated directly from the data, but some need to be swept, matching an output of the model to the corresponding value extracted from data. Moreover, we need different functional relationships coming directly from data.

3.3.1 Functional relationships

The first functional relationship we use is the one between the traffic and the delay. In order to compute it for each airport, we used DDR data, with the use of the difference between M3 (actual, radar-tracked flight) and M1 (last-filed flight plan) departure times as an approximation for the real delay. In appendix B.1.1, we study the potential difference between these delays and the real ones, that are only available at an aggregated level from CODA data (and not throughout the day). We use one month of data, and we compute the number of departures during each hour of each day, as well as the corresponding mean delay. We then average them over the whole month of data and for each airport, and perform an exponential fit. There are only two parameters to the fit, one of them being a measure of the capacity. Finally, the airports were gathered according to the three clusters found in Section 3.1.2.3, which means that we use only three relationships in total.

The passenger costs cover compensation and duty of care, as required by Regulation 261 [53], and also reduced market share costs arising from poorer punctuality (they do not include (internalised) passenger value of time costs). The delay costs are sourced from [3]. Based on this source, we carried out a regression fit for *primary* delay (to avoid double-counting across the network by including reactionary impacts) costs using the weights of the aircraft and the delay durations. The final function is:

$$c_d = -7.0 \delta t - 0.18 \delta t^2 + (6.0 \delta t + 0.092 \delta t^2) \sqrt{MTOW}, \quad (3.1)$$

For the model, we set \sqrt{MTOW} to its average across all aircraft departing from the airport.

In this equation, the delay should not be an average one, but the actual delay of each flight. This is a potential issue for the model, given that we are using average delays (per hour). Indeed, it is clear, for instance, that a null average delay still produces a cost, since some flights are still delayed, and thus bear a cost, whereas airlines with flights ahead of schedule are assumed not to benefit financially (they may even suffer costs). Moreover, the cost is super-linear with the delay, which means that long delays cost proportionally much more than shorter ones. As a consequences, we used DDR data to compute the **intra-hour** distribution of delays. These distributions are fitted with log-normal distributions (see Annex B.1.3), which is a suitable approximation, even though very long delays are sometimes underestimated. The previous equation is then corrected, using the **expected** values of the cost, based on the probability density of delays. This procedure is undertaken airport by airport, increasing the total real cost. Note that, in particular, the expected cost of delay for a null average delay is not null any more.

3.3.2 Direct calibration of parameters

Some parameters can be directly estimated from the data, for each airport:

- The load factor l_f is given by the ratio of the number of flights and the number of passengers (this is not the real average load factor which would need to consider each flight separately).
- The (average) aeronautical revenues per flight P are given by the total aeronautical revenues divided by the number of flights.
- The (average) non-aeronautical revenues per passenger w are given by the total non-aeronautical revenues divided by the number of passengers.
- The distribution of traffic $\{T\}$ through the day is fixed by averaging one month of data, splitting the day into 24 hour periods.

The value of time for all passengers, not useful *per se* for the model but impacting on passenger satisfaction, can be found in the literature. To have a more realistic description, we decided to use two values of time, which are usually associated with business (v_b) and leisure passengers (v_l), sourced from [55]. We then consider that most passengers on low-cost carriers have a lower value of time – often associated with leisure trips – whereas passengers travelling with traditional/full service airlines have a higher value of time, reflecting more business trips. As a consequence, the average value of time in our model is:

$$v = v_l r_{lcc} + v_b (1 - r_{lcc}),$$

where r_{lcc} is the share of low-cost carriers at the airport. This value is also directly taken from data (i.e. DDR data). Note that this is a very crude approximation, since many business-purpose trips are made on low-cost carriers, for example. However, since precise passenger profiles are not currently available to the research team, this is a reasonable approximation, and better than simply considering as equal business and leisure passenger volumes across all flights.

Finally, an important parameter is the marginal cost of extra capacity per passenger α . This value is quite difficult to extract from data, so we consider it as a free parameter most of the time. However, in order to have an approximate idea of its value, one can assume that the main objective of an airport is to deliver capacity, and thus all of its costs should be related to this delivery. A regression of the total costs is thus carried out (outlay on infrastructure, such as a new runway) for the various airports, as a function of capacity. A linear regression is reasonable in this case ($R^2 = 0.53$) except for the high cost airports (essentially CDG, LHR, and FRA). The cost is thus around 24 000 euros per unit of capacity per day. This corresponds to 220 euros per passenger, per day, taking CDG as an example.

3.3.3 Post-calibration of parameters

We call post-calibration the operation of running the model with different values of one or more parameters, and comparing an output of the model to some values extracted from the data.

With the final version of the model, we only need to post-calibrate one parameter (β), tuning the demand at the airport. More specifically, it is a multiplicative factor for the traffic. When calibrating this, we match the effective number of flights operated at the airport with the real ones in the data, sweeping the value of β . After calibration, β is usually a value greater than 1. Indeed, we need to increase the traffic volumes obtained from data, since the probability that the airline operates the flight is smaller than 1. In other words, we compensate for the disincentive of airlines (i.e. due to the cost of delay) associated with an increase in demand, in order to match the observed number of flights.

3.3.3.1 Summary of calibration

In summary, the calibration includes the following steps:

- Maximum take-off weight $MTOW$ is included in the cost-delay relationship.
- Average load factor l_f , aeronautical revenues per flight P and non-aeronautical revenues per passenger w , value of time v , total initial cost c_{init} , and distribution of traffic $\{T\}$ through the day are taken directly from data.
- Fitting parameters cc and \mathcal{C}_{init} (the latter being the capacity) for delay-traffic load relationship are set.
- Cost of delay relationship is corrected based on intra-hour log-normal fitting distributions of delays.
- Demand factor β is post-calibrated by matching of the number of flights with the data.

The ‘total initial cost c_{init} ’ represents the total current costs of the airport, i.e. costs for providing the current capacity. Finally, we have two parameters remaining, the smoothness of the airline decision s and the marginal cost of capacity α , which is the operational cost of running one extra unit of capacity. The latter could be estimated, for instance, with the current cost per passenger, but we prefer to keep it as a variable, since the exact figure is likely to be different from just ‘capacity/passengers’.

The smoothness is thus the last free parameter of the model. It represents the sensitivity of the airline to the cost of delay. The higher it is, the smoother the decision will be, i.e. the airline will not *suddenly* cease operating at an airport when

the cost increases. When the decision is harsh (low value of s), the airline ceases to operate the flight as soon as the cost is non-zero.

It is clearly linked to the elasticity of demand (associated with the airline) and the cost (of the airline operating a flight at the airport), which is very hard to estimate because of lack of detailed airline data. It is worth noting, however, that:

- a basic sensitivity analysis (see Annex C) shows that the results of the model do not depend strongly on the value of s ,
- the parameter is actually not totally free, but is constrained at low values. This comes from the fact that a low elasticity cannot fulfil demand requirements.

As a result, the model is only slightly over-fitted and thus suitable for the task.

In Table 3.4, we present a summary of all the parameters¹ and their types, corresponding to the ways that they are calibrated.

Name of parameter	Short description	Type of parameter	Value for calibrated airport
$MTOW$	Max. take-off weight	DC	120 euros
l_f	Load factor	DC	271 pax/flight
P	Airport charges	DC	2393 euros
C_{init}	(Departure) capacity	DC	35.59 flights
cc	Delay at zero traffic	DC	-1.83 minutes
v	Value of time	DC	39.91 euros
T	Distribution of traffic	DC	(distribution)
w	Average revenue per passenger	DC	17.78 euros
c_{init}	total initial cost	DC	3.3m euros
β	Traffic multiplier (demand)	PC	3.45 (no units)
α	Marginal cost of capacity	FP	–
s	Smoothness	FP	–

Table 3.4: List of parameters of the model, with their types related to calibration. DC: Direct calibration, FP: Free parameter, PC: Post-calibrated. The last column presents the value obtained for the calibration with the selected airport.

Once the model is calibrated, it is possible to change some of the parameters to see the impact on other variables, as is demonstrated in the next section.

¹This list of parameters is different from the metrics considered in the data analysis section. Indeed, only those parameters were selected that could be included in the model and calibrated. Many of the parameters of the data analysis were included at first in the model, but were then removed as progress was made to a fully calibrated model.

4 Model implementation and results

4.1 Implementation

Summary

The model is implemented in Python and uses a MATLAB graphical user interface (GUI) for user-friendly interactions. The user can change the values of different parameters, and the output representation, for a seamless experience of the model.

4.1.1 Engine implementation

The model is written in Python and is divided into two parts. The first relates to the calibration, where relevant data are extracted from our sources and several operations are performed to prepare the model, such as regressions. The output of this part is a Python object embedding all the calibrated parameters. The second part of the code is the object itself, which is able to produce different outputs based on the free variables passed in the inputs. The output is then used by the graphical user interface (GUI), described below.

4.1.2 Graphical user interface (GUI)

The model developed is aimed at experts in air transport, that do not necessarily have the technical or programming skills to execute or modify a software platform. In order to make the model more accessible, a visualisation layer, or GUI, has been developed on top of the data-driven, ‘back-end’ model (engine). The model is then delivered as an autonomous piece of software usable without programming skills.

This visualisation helps the user to understand the underlying model behaviour and evolution when varying certain input parameters and airport types, for example determining the combination of parameters that lead to desirable outputs or, in some

cases, optimum values. The visualisation tool also helps to determine the stability and sensitivity of the optimal points, the local behaviour in small neighbourhoods, visually.




Figure 4.1: Screenshot of the visualisation layer.

The engine has been developed in MATLAB and can be deployed on Windows, Mac and Linux platforms. It is packed as a MATLAB standalone application. Therefore, a licensed copy of MATLAB is required. If not available, MATLAB runtime needs to be installed. (Instructions for installing MATLAB standalone applications and runtime libraries can be found at [56].) The engine is compatible with modules (airport and airline models) written either in MATLAB or Python programming languages and exports output data into common formats: .png for figures, plus .xml and Excel-compatible .csv files for tables. Although the visualisation layer has been developed as intuitively as possible, embedded in the software, there is a user manual to facilitate the use of the GUI. It can be accessed from the ‘Help’ menu at the top right-hand corner of the application window, as shown in Figure 4.2.

Source codes are stored in a private and secure GitHub repository. GitHub is an online control version repository with collaborative capabilities. Updates, bug fixes and new versions of the model will be released beyond the current project. Snapshots as a MATLAB standalone application could be made on demand. The source code will be hosted and maintained indefinitely.




Figure 4.2: Help menu.

4.2 Calibrated model results

Summary

The model calibrated on a large, European hub airport shows the presence of an optimum in capacity greater than the current capacity, for a wide range of marginal operational costs. At the optimum, the mean delay has fallen by approximatively 5% and the passenger utility has increased by the same amount. When calibrating different airports, one can compare the maximum operational cost for which a given increase is profitable to the airport. As expected, bigger airports can, in general, sustain greater extra operational cost, in particular because they are the most congested. However, the number of passengers is not the only factor, and some large airports cannot sustain high operational costs.

4.2.1 Calibrated airport

We have two free variables. The first is the new capacity \mathcal{C} targeted by the airport, and the second is the marginal cost of capacity α . All other variables have been calibrated, as described previously.

Figure 4.3 shows the net income of the airport when changing these two variables. The global maximum is clearly reached when the marginal cost is null and the capacity is at its maximum value. It is interesting that, for a given value of α , there are two possibilities: either there is a maximum in capacity, or there is not. Indeed, for a small value of α , when the capacity increases, the net income also increases, up to a certain point, after which it starts decreasing. This is because increasing the capacity

decreases the delay, and more flights are operated. However, in contrast, the income per passenger decreases with capacity, as we explain below.




Figure 4.3: Evolution of the net income of the airport as a function of the capacity and the marginal operational cost.

In order to see this effect more clearly, we present various graphs in Figure 4.4, where we have fixed $\alpha = 60k$. Starting from the top left, graph (a), we first show the revenues per passenger, which in our model is a constant. Indeed, only the number of flights, and thus the number of passengers, change when the capacity is altered. Note that in reality, there might be two different opposite effects related to this. On the one hand, the non-aeronautical revenues have been found to be weakly correlated with passenger satisfaction. As a result, increasing capacity may increase the passenger satisfaction, and hence the non-aeronautical revenues per passenger. On the other hand, it may be possible that some diseconomies of scale might occur when the airport is handling more passengers and reaches a certain size threshold, as discussed in the literature review.

Graph (b) shows the cost of capacity, which increases with the number of passengers. This is the case because increasing capacity by one unit does not bring all the potential passengers, but only a fraction of them. As a result, the income per passenger decreases with the capacity, as shown in graph (c).

In graph (d), we show that average delay decreases with capacity, which is the main objective. As a consequence, the cost of delay for airlines decreases too (graph

(e)), and the airport serves the demand better (graph (g)), with a higher number of flights (graph (h)).

The relationships between the increase in the number of passengers and the decrease in the net income per passenger leads to a maximum for total net income (at least for some values of α), see graph (i)). Finally, as average delay decreases, passenger utility monotonically increases (i.e. with better satisfaction), see graph (j)).

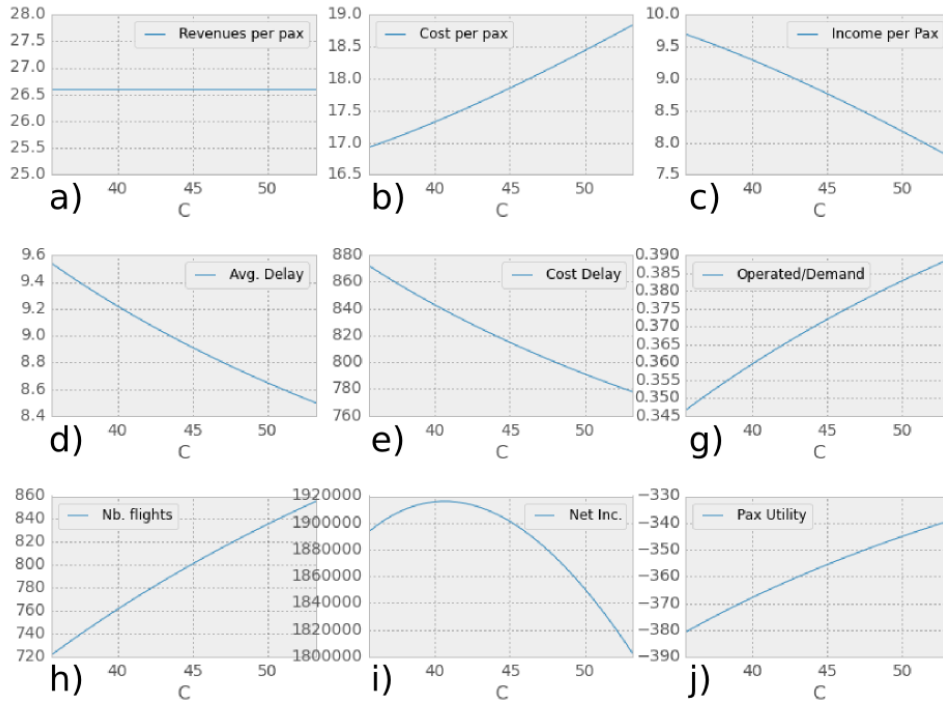


Figure 4.4: All results are expressed per day. Evolution of (a) revenues per passenger, (b) cost per passenger, (c) net income per passenger, (d) average delay per flight, (e) cost of delay per flight, (g) ratio between operated flights and demand, (h) number of flights actually operated, (i) total net income of the airport, and (j) passenger utility, as functions of the airport capacity.

4.2.2 Comparing different airports

In order to compare the different airports, we identify the value of α for which the net income of the airport reaches a maximum, when the capacity is increased by 50%

with respect to its current value. This gives a rough idea of the maximum marginal cost at which it is profitable to increase the capacity. Table 4.1 show the results for various airports, as well as some other metrics, for comparison. It is striking to see how this typical marginal cost varies significantly among airports, from LPPR with 541 euros, to EGLL with 83 540 euros. This means, roughly, that an increase of 50% of the capacity at Heathrow is profitable up to a marginal operational cost of 83 540 euros per day per unit of capacity. On the other hand, an increase of 50% of the capacity for Porto airport would be profitable only if the marginal operational cost were below 541 euros per day.

	C	α	Pax (millions)	CUI
EDDF	37.6	43087	59.6	0.54
EDDH	14.3	10404	14.8	0.49
EDDK	40.7	5106	9.45	0.53
EDDL	35.5	5106	21.9	0.49
EDDM	73.5	19038	39.7	0.52
EFHK	29.6	5586	15.9	0.44
EGBB	10.5	1424	9.72	0.45
EGCC	14.9	3409	22.1	0.47
EGKK	20.2	6127	38.1	0.61
EGLL	28.7	83540	73.4	0.59
EGSS	17.5	7148	19.9	0.49
EHAM	50.2	29615	55.0	0.53
EIDW	45.5	4085	21.7	0.59
EKCH	45.0	8477	25.6	0.49
ESSA	30.8	3179	22.5	0.53
LEMD	29.9	17687	41.8	0.55
LFPG	41.3	53073	63.8	0.54
LGAV	13.1	7014	15.2	0.44
LIRF	288.9	8323	38.5	0.51
LOWW	34.3	8784	22.5	0.46
LPPR	4.7	541	6.93	0.37
LPPT	16.5	2234	18.1	0.49
LSZH	26.1	8630	25.4	0.53
LTBA	12.8	3179	56.7	0.67

Table 4.1: Capacity, typical marginal cost for which an increase of 50% of the capacity is profitable, number of passengers and capacity utilisation index for all calibrated airports.

It is interesting to see that this threshold seems to be highly correlated with the

size of the airport (see Figure 4.5), in a quite non-linear way (the initial analysis shows that it is compatible with a power law). This could be related to economies of scale as discussed in the literature review.




Figure 4.5: Threshold of the marginal cost of capacity for which an increase 50% of the capacity is profitable for different airports.

4.2.3 Effect of parameters on the position of the optimum

In this section we examine the effect of the various parameters on the position of the optimum. Since all the parameters are calibrated, or considered as variables, some parts of the calibration procedure are relaxed to allow some of the parameters to vary.

The first parameter that we are interested in is the load factor. As highlighted in the literature review, the average load factor at airports has been increasing in recent years, except at highly congested airports. The main reason is that this is a relatively cheap way of increasing their effective capacity. As a consequence, many of them try to find incentives for airlines to increase the size of their aircraft.

In the calibrated model, the load factor l_f is derived simply by dividing the total number of passengers by the total number of flights for a given airport. Here, we relax this calibration and sweep the parameter, every other parameters being unchanged. This load factor in our model has a very simple effect: it increases the non-aeronautical revenues linearly. As a consequence, the optimal capacity is shifted, as shown in Figure 4.6. The figure shows the optimal capacity as a function of the load factor.




Figure 4.6: Optimal capacity as a function of the average load factor.

For low load factors, the optimal capacity is at the current capacity or below (by design, here the capacity cannot fall below the current value). At a certain point, the load factor is high enough, and the optimal capacity starts to increase linearly with it. This simple effect is due to the fact that higher load factors act as a “free” increase in capacity, and thus the airport has less need to increase actual infrastructure.

The second interesting feature that we study is the effect of **predictability** when **punctuality** is fixed. In other words, the effect of the variance of the delay when the mean delay is fixed. To simulate this effect, we change the ‘correction term’ of the cost of delay function, reducing progressively the variance of the log-normal distribution while keeping a constant mean.

First in Figure 4.7, we show the impact of decreasing predictability on the calibrated model with the current capacity, for different metrics. The right-hand side of each graph corresponds to the situation in the data, i.e. a normalised standard deviation of 1. The extreme left-hand side corresponds to a perfectly predictable situation, but not perfectly punctual. This last limit is not very realistic and the model probably fails in this region, due to the fact that in reality airlines would integrate a simple off-set of departure times into their objective functions.

Nevertheless, moderate increases of predictability (decreases of standard deviation) are probably well captured by the model. Unsurprisingly, the net income starts by increasing when the standard deviation falls. This is due to the fact that it is less costly for airlines to operate at this airport, and thus they are more willing to do so. However, the net income is not monotonic - we return to this point later.

We first focus on the average delay at the airport. The counter-intuitive results

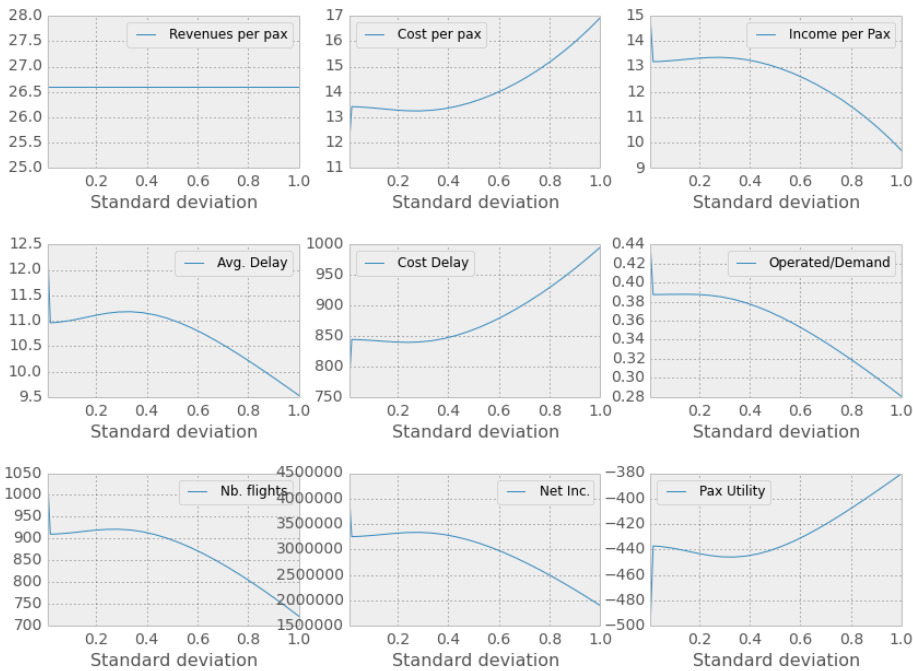


Figure 4.7: Evolution of different metrics with the increase of the standard deviation of delay at the airport (decrease of predictability).

from the model indicate that increasing the predictability actually increases the mean delay at the airport. Apparently, the airport trades the extra predictability against some extra average delay, in some kind of mean-variance trade-off. In other words, the airport can have a higher mean delay and still be more attractive to the airlines because of the increase in predictability. Associated implications, such as those related to strategic buffers, could be modelled in future.

Driving the standard deviation even lower, the airport reaches a maximum in its net income, before this decreases slightly and then increases substantially. The results for the extreme values are probably not reliable, but the moderately low values could still be meaningful. To understand the latter properly, further study would be required.




Figure 4.8: Optimal capacity as a function of the standard deviation of delay.

The consequences of this complex behaviour on the position of the optimal capacity are quite difficult to predict. Indeed, as shown in Figure 4.8, the optimal capacity does not vary monotonically with the predictability. When the latter increases, the optimal capacity falls, reaching a minimum higher than the current capacity, before increasing again. This is perfectly consistent with the fact that average delay first increases with the predictability in Figure 4.7. Indeed, increasing predictability can be seen here as an effective reduction of capacity.

These results shed light on a *potentially* deleterious effect relating to a reduction of predictability, namely that the airport might react in such a way that punctuality decreases.

4.3 Exploratory results

Summary

The fully calibrated model has a very simple, constant non-aeronautical revenue per passenger type. In this section, this assumption is relaxed and more complex behaviours are explored. In particular, passengers can now shop more when they are delayed, but since their satisfaction decreases, they might spend less per unit of time. Since there are no data supporting the calibration of this part, the potential impacts are illustrated by choosing some *ad hoc* functions. The results show that several optima in capacity can occur for some values of the parameters. In particular, a small maximum can appear at low values of capacity, whereas a higher one (more profitable) occurs at larger capacity increases. This means that a risk-averse airport could decide to reach only the first maximum instead of the second one, thus not maximising its profit - leaving passengers and flights with higher delays.

In this section, we slightly modify the model in order to explore its possibilities. These results are less realistic, because they rely on more assumptions. In particular, there are more free parameters, which are impossible to calibrate with the current data. However, we include this part because it allows us to explore the missing features of the model and their potential consequences.

The general feature we add to the model is the possibility of modifying the non-aeronautical revenues per passenger. In particular, we decide to allow the passengers to spend more at the airport: (1) if their flight is delayed; (2) if their overall satisfaction increases. These two effects are opposite to each other (because the general satisfaction should decrease when the delay increases), but it is quite likely that they do not cancel each other out. Indeed, we assume here that with short delays the passengers spend more at the airport, whereas satisfaction is likely to fall with longer delays. In order to model this, we choose a linear dependence for the first effect, and a quadratic one for the second. Hence, the typical revenues per passenger will look like the curve plotted in Annex B.3, in Figure B.5. The exact equation we used is also shown in the annex. In the following, we show some examples of what could happen with this effect, using the calibrated airport.

4.3.1 ‘Better shopping time’

It is first interesting to see what happens with only one effect or the other. In Figure 4.9, we show what happens when we include only the effect of changing satisfaction levels. As expected, the revenue per passenger is now increasing with the capacity of the airport (as passengers are more satisfied). This has the consequence of creating a net income per passenger that first increases and then decreases, which is contrary to the results obtained for the fully calibrated model. However, the effect on the total net income of the airport is qualitatively null. If the maximum is clearly shifted towards higher capacity, resulting in higher net income, there is still only one maximum, as previously highlighted for the fully calibrated model.

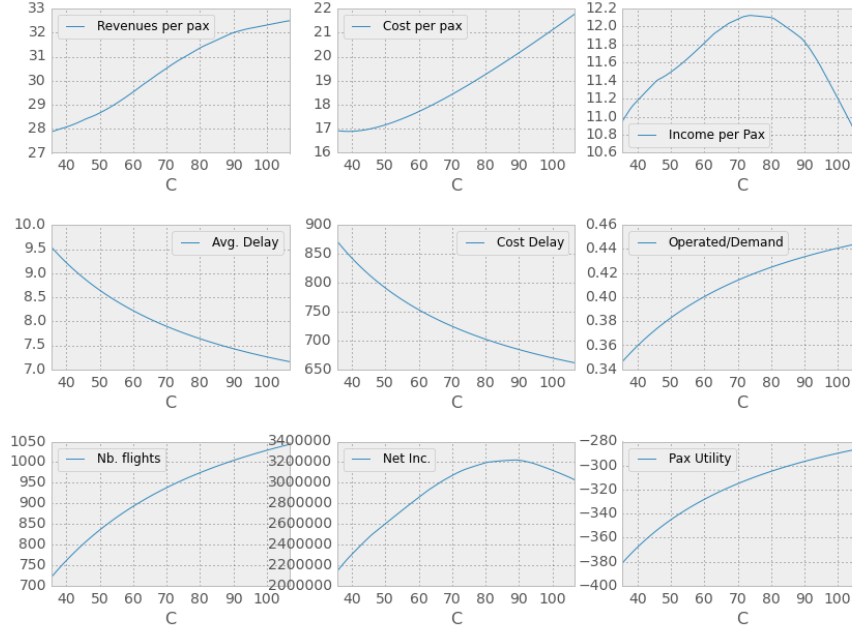


Figure 4.9: Different outputs of the models with satisfaction effect (see caption of Figure 4.4 for the meaning of each graph). These plots have been obtained with the following values of free parameters: $t_e = 0$, $s_e = 1500$, $cap = 24$, $\alpha = 40000$.

4.3.2 ‘Longer shopping time’

When disabling the effect of satisfaction and keeping only the ‘more shopping time’ effect, the results for net income are again quite similar (not shown here). Indeed, the

revenue per passenger is now clearly decreasing with capacity, and since the cost is increasing, the net income per passenger also decreases monotonically. As a result of the increase in the number of flights, there is still a maximum for the net income, which is clearly shifted towards the low capacities this time. Note that this is interesting *per se*, because in this case the airport has much less incentive to increase capacity.

4.3.3 Both effects

Finally, we run the two effects at the same time. Figure 4.10 shows the same type of plots as before, which are now more complex. First, the revenues per passengers are no longer monotonically decreasing. At first, the increase in capacity shortens the shopping time and decreases the revenues, but then the increase in satisfaction takes effect and the revenues increase again. The effect on the income per passenger is dampened by the costs, and thus still monotonically decreasing. However, the relationship between the revenues per passenger and the number of passengers, creates an interesting pattern for the total net income. Indeed, this first increases slightly because of the rise in passengers, which is strong at the beginning, then decreases because the revenues per passenger fall. It then increases again when the satisfaction effect exerts itself, and finally drops for high capacities: being significantly affected by the cost of capacity. This is an interesting pattern, because it means, for instance, that the airport could be trapped in the first, local, maximum without sufficient incentive to reach higher capacity, which is much more expensive in terms of investment.

The important conclusion of this section is that the modification of the revenues per passenger, which is clearly linked to the economic efficiency of the airport, can blur the picture drawn by the fully calibrated model and create some local maxima for the expected net income of the airport. More specific studies would be needed to discover if one effect or the other – or both – is more likely to happen in reality. It is important to note that aeronautical revenues are not considered here because of the complexities identified in the literature review.

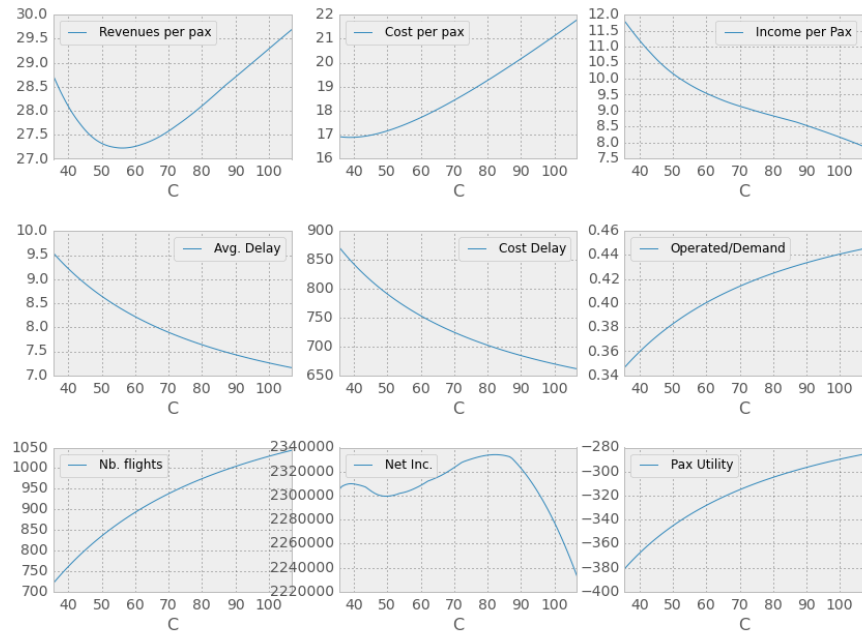


Figure 4.10: Different outputs of the models with both effects (see caption of Figure 4.4 for the meaning of each graph). These plots have been obtained with the following values of free parameters: $t_e = 20$, $s_e = 1500$, $cap = 24$, $\alpha = 40000$.

5 Conclusions

The work presented in this deliverable aimed at producing a new, highly data-driven model able to give quantitative insights into the costs and benefits of extra capacity at an airport.

Technically, one of the most difficult tasks was the acquisition of data, and their consolidation. Indeed, it took a large amount of the effort just to acquire, clean, extrapolate and consolidate the data. Most of the sources have been used one way or another, but two severe limitations have hindered the modelling process.

First, the lack of times series is a problem when one tries to build causal relationships, because correlations between different airports are typically less interpretable than correlations in time for a single airport. Second, we were able to calibrate the model only for a few airports, because there were several fields missing in the data for most of them. An analysis of more airports, especially for the cluster analysis, would have allowed us to differentiate the model better, based on different types of airport.

A further main challenge was producing a model which can be calibrated, but with enough realism to produce some quantitative outputs. Several tests have been made (we have produced around 15 versions of the model in total), but the final one was simplified as much as possible to allow a full calibration: only one free parameter remains, which can be in fact roughly estimated, as shown in the sensitivity analysis [C](#).

However, if the model is functionally simple, it embeds potentially complex behaviour through the use of full distributions of delay and the corresponding expected costs. In fact, special attention has been applied to the impact of the non-linearity of the cost of delay, which can be taken into account only by using the full distribution of delays (in our case, a log-normal fit for each airport and each hour of the day). For the same reason, it was important for the model not to take for granted the basic definition of airport capacity, but to use a data-driven relationship, which, in essence, includes uncertainties regarding delays – with the consequence, for instance, that the cost of delay is not null even when the airport is operating below its theoretical capacity threshold.

Another advantage of a functional model is that it is very easy to modify and to include different relationships, either for further exploratory research or for more airport-tailored results. This was performed in the more exploratory part of the results, where we showed that complex interactions may arise between time spent shopping, passenger satisfaction, cost of capacity and satisfying demand – creating different types of curves for net income. This, in turn, will likely influence the future choices of the airport with respect to potential extensions of capacity.

However, it is clear that a functional model such as this one cannot capture all the complexity of reality. For instance, compared to some other simulation-based models, it is hard to introduce the feedback between different airports. Indeed, it is clear that airlines make decisions at *least* on a two-airport basis (departure and arrival), and, in some cases, model route impacts at the full network level.

Hence, airlines will react differently to an increase in capacity at a departure airport depending on the current state (and future, expected state) of the arrival airport. An increase or decrease in delay will thus meet different reactions, based on the delay at the other airport.

Another point not captured by our current model is the diversity of airlines. Not all of them react in the same way to the cost of delay, because they have different business models, different networks, different schedule buffers, etc. An operator such as Ryanair, for instance, known to deploy large schedule buffers, may be more flexible in terms of delay than some other operators, e.g. those operating connecting wave structures at given airports (especially hubs).

The business models of the airports themselves will also attract different types of operator, and increases in capacity will have different impacts on them. Stated more generally, the heterogeneity of the different agents cannot be well captured by this type of representative agents model.

Particular care has been taken regarding the choice of the right metrics and mechanisms to consider in the model, constrained by the data available. For example, delay at an airport can obviously have many causes. Since it was desirable to capture the delay directly caused by congestion, turnaround delay was focused upon. It was also important to treat aeronautical and non-aeronautical revenues separately, since they are two distinct sources that the airport can manipulate (to varying extents) to maximise revenue. In particular, as seen in the exploratory analyses, airports may have a strong incentive to increase delay in some cases.

We have demonstrated that the model can be easily calibrated on real data and runs very well for the airports in our dataset. It produces reliable and realistic results.

Some of them, in particular, are worth noting.

First, all the airports exhibit a maximum in net income as a function of capacity, if the marginal cost of operating extra capacity is sufficiently low (as illustrated in Figure 5.1). This threshold in the marginal cost is, however, rather different across airports, and only a few airports can sustain a high cost of capacity. As expected, these are the largest and most congested airports, which clearly need extra capacity. Moreover, we have shown that this threshold is roughly consistent with the airports' current operational cost of capacity, which means that they should be able to manage this growth, subject to the availability of investment.




Figure 5.1: Illustration of optimum capacity at an airport.

The exploratory results show that the picture can be significantly modified by the introduction of variable, non-aeronautical revenues per passenger as illustrated in Figure 5.2. When tendencies to ‘shop more with more time’ and ‘shop better with increased satisfaction’ are introduced, the net income can exhibit different maxima and minima. The direct consequence is that an airport would probably not be willing (or able) to invest sufficient capital to reach the global maximum, and is likely to be ‘trapped’ in a local maximum. Since an increase in capacity is incremental (e.g. new runway, new terminal), this may actually render it impractical for the airport to reach any maximum.

In future work, we hope to carry out a full calibration exercise on a specific airport.




Figure 5.2: Complex landscape with local minima arising from the non-trivial non-aeronautical revenues per passenger.

A Data analysis

In this annex, we give some detail about the data analysis performed in support of the modelling process.

A.1 Correlation structure

The full correlation matrix can be found in the next three pages. Table A.1 shows the meaning of each of the variables:

Abbreviation	Meaning
RegionE	Share of European flights
Airlines	Number of air operators
Routes	Number of routes
Terminals	Number of terminals
Gates	Number of gates
Runway number	Number of runways
Total pax	Total number of passengers
Delay per flight	CODA turn-around delay per flight
Avg pax/route	Number of passengers per route
Net basic utility	Net income minus aggregated cost of delay
Runway usage	Number of flight per runway
Terminal usage	Number of flight per terminal
Avg pax/flight	Number of passengers per flight (load factor)
FLT_DEP_1	Number of flights
turn_delay	Total turn-around delay
Capacity (hour)	Maximum number of movements per hour
CUI	Capacity utilisation Index
Peak hour vol	Peak hour volume
Dom pax 2014	Number of domestic passengers
Int pax 2014	Number of international passengers
Public ownership	Share of public share in airport
CHT%	Percentage of charter flights
FSC%	Percentage of full-service flights
LCC%	Percentage of low-cost flights
REG%	Percentage of regional flights
OTH%	Percentage of other flights
TotalAeroRev13	Aeronautical Revenues
TotalNoneAeroRev13	Non-aeronautical revenues
TotalOpExp13	Operational expenses
NetOpIncome13	Operational net income
Satisfaction	Degree of passenger satisfaction at the airport
Coordination	Degree of coordination of the airport
Low	Aggregated market share of regional and scheduled flights
High	Aggregated market share of low-cost and charter flights

Table A.1

	RegionE	Airlines	Routes	Terminals	Gates	Runways Number	Total pax	Delay per flight	Avg pax/ route	Net basic utility	Runway usage	Terminal usage	Avg pax/ flight
RegionE	1.00												
Airlines	-0.67	1.00											
Routes	-0.74	0.82	1.00										
Terminals	-0.39	0.54	0.46	1.00									
Gates	-0.65	0.74	0.78	0.32	1.00								
Runway number	-0.14	0.29	0.21	0.28	0.18	1.00							
Total PAX	-0.72	0.80	0.90	0.49	0.78	0.25	1.00						
Delay per flight	-0.52	0.53	0.63	0.34	0.38	0.20	0.60	1.00					
Avg pax/route	-0.02	0.12	-0.06	0.16	0.17	0.07	0.16	0.11	1.00				
Net basic utility	0.17	-0.05	-0.42	-0.15	-0.27	-0.28	-0.36	-0.33	-0.17	1.00			
Runway usage	-0.61	0.41	0.64	0.10	0.58	-0.25	0.67	0.36	0.28	-0.14	1.00		
Terminal usage	-0.61	0.63	0.78	0.02	0.69	0.25	0.78	0.31	0.19	-0.27	0.50	1.00	
Avg pax/flight	-0.45	0.34	0.52	0.25	0.35	0.12	0.61	0.42	0.56	-0.10	0.35	0.43	1.00
FLT_DEP_1	-0.77	0.74	0.91	0.36	0.73	0.43	0.98	0.58	0.45	-0.42	0.68	0.68	0.48
turn_delay	-0.72	0.71	0.83	0.45	0.58	0.41	0.91	0.80	0.37	-0.45	0.56	0.53	0.47
Capacity (Hour)	-0.71	0.84	0.90	0.49	0.80	0.29	0.94	0.59	0.17	-0.49	0.63	0.79	0.45
CUI	-0.59	0.70	0.77	0.38	0.69	0.16	0.75	0.47	0.21	-0.19	0.63	0.70	0.54
Peak hour vol	-0.72	0.85	0.90	0.49	0.81	0.29	0.94	0.58	0.17	-0.49	0.63	0.79	0.45
Dom pax 2014	-0.56	0.52	0.61	0.27	0.52	0.14	0.69	0.10	0.11	-0.21	0.38	0.64	0.44
Int pax 2014	-0.74	0.79	0.88	0.48	0.76	0.25	0.98	0.63	0.14	-0.35	0.63	0.76	0.56
Transfer pax 2014	-0.66	0.54	0.74	0.38	0.62	0.42	0.90	0.50	0.66	-0.55	0.50	0.64	0.49
Public ownership	0.19	-0.10	-0.11	0.11	-0.03	0.12	-0.10	-0.22	0.06	-0.25	-0.12	-0.10	-0.16
CHT%	0.20	-0.09	-0.18	-0.06	-0.16	-0.08	-0.12	-0.21	0.43	0.07	-0.10	-0.12	0.09
FSC%	-0.42	0.45	0.36	0.28	0.42	0.18	0.42	0.39	0.05	-0.26	0.30	0.37	0.10
LCC%	0.13	-0.17	-0.01	-0.09	-0.10	-0.08	-0.10	-0.08	-0.08	0.21	-0.06	-0.03	0.34
REG%	0.15	-0.15	-0.18	-0.13	-0.16	-0.05	-0.17	-0.20	-0.14	0.08	-0.20	-0.20	-0.37
OTH%	0.20	-0.29	-0.30	-0.17	-0.29	-0.06	-0.30	-0.38	-0.14	0.09	-0.29	-0.29	-0.58
TotalAeroRev13	-0.67	0.58	0.58	0.37	0.58	0.24	0.80	0.47	0.76	-0.01	0.58	0.37	0.53
TotalNonAeroRev13	-0.64	0.56	0.72	0.39	0.52	0.36	0.85	0.57	0.67	-0.47	0.51	0.49	0.45
TotalOpExp13	-0.65	0.53	0.72	0.40	0.59	0.34	0.87	0.57	0.71	-0.57	0.55	0.49	0.46
NetOpIncome13	-0.61	0.57	0.50	0.34	0.45	0.20	0.71	0.40	0.67	0.22	0.53	0.33	0.47
Satisfaction	-0.45	0.36	0.54	-0.02	0.50	0.24	0.53	0.30	0.48	-0.10	0.43	0.53	0.42
Coordination	-0.18	0.51	0.49	0.28	0.41	0.18	0.42	0.44	0.07	-0.18	0.18	0.36	0.10
Low	0.21	-0.21	-0.09	-0.11	-0.17	-0.12	-0.15	-0.14	0.12	0.21	-0.09	-0.08	0.34
High	-0.30	0.33	0.23	0.18	0.29	0.14	0.28	0.28	-0.05	-0.23	0.20	0.21	-0.07

Table A.3: Pearson correlation coefficients.

	FLT_DEP_1	turn_delay	Capacity (Hour	CUI	Peak hour vol	dom pax 2014	int pax 2014	transfer pax 2014	Public Ownership	CHT%	FSC%	LCC%	REG%	OTH%
RegionE														
Airlines														
Routes														
Terminals														
Gates														
Runway number														
Total pax														
Delay per flight														
Avg pax/route														
Net basic utility														
Runway usage														
Terminal usage														
Avg pax/flight														
FLT_DEP_1	1.00													
turn_delay	0.91	1.00												
Capacity (Hour)	0.97	0.89	1.00											
CUI	0.78	0.66	0.79	1.00										
Peak hour vol	0.97	0.88	0.99	0.77	1.00									
Dom pax 2014	0.54	0.36	0.64	0.59	0.66	1.00								
Int pax 2014	0.95	0.93	0.92	0.70	0.91	0.53	1.00							
Transfer pax 2014	0.90	0.85	0.83	0.55	0.82	0.39	0.90	1.00						
Public ownership	-0.04	-0.10	-0.07	-0.09	-0.07	-0.06	-0.12	-0.10	1.00					
CHT%	-0.22	-0.21	-0.13	-0.07	-0.13	-0.16	-0.11	-0.16	-0.01	1.00				
FSC%	0.63	0.57	0.46	0.34	0.47	0.33	0.41	0.58	-0.07	-0.21	1.00			
LCC%	-0.33	-0.29	-0.14	-0.05	-0.15	-0.05	-0.10	-0.31	-0.06	-0.14	-0.59	1.00		
REG%	-0.22	-0.22	-0.16	-0.18	-0.15	-0.07	-0.20	-0.24	0.10	-0.10	-0.29	-0.35	1.00	
OTH%	-0.40	-0.36	-0.29	-0.27	-0.30	-0.29	-0.28	-0.27	0.13	-0.09	-0.26	-0.16	-0.04	1.00
TotalAeroRev13	0.73	0.73	0.67	0.56	0.67	0.31	0.83	0.74	-0.10	-0.17	0.40	-0.23	-0.10	-0.36
TotalNonAeroRev13	0.84	0.89	0.80	0.54	0.80	0.29	0.89	0.90	-0.03	-0.21	0.51	-0.35	-0.16	-0.30
TotalOpExp13	0.85	0.88	0.83	0.56	0.83	0.34	0.90	0.93	0.03	-0.19	0.48	-0.31	-0.15	-0.33
NetOpIncome13	0.64	0.65	0.56	0.48	0.56	0.20	0.76	0.62	-0.16	-0.16	0.38	-0.24	-0.10	-0.29
Satisfaction	0.56	0.44	0.55	0.37	0.59	0.13	0.56	0.52	-0.24	-0.17	0.30	-0.06	-0.16	-0.35
Coordination	0.26	0.32	0.48	0.43	0.48	0.20	0.40	0.29	0.01	-0.03	0.08	0.04	0.02	-0.23
Low	-0.37	-0.33	-0.20	-0.08	-0.20	-0.13	-0.14	-0.34	-0.06	0.32	-0.66	0.89	-0.38	-0.19
High	0.51	0.46	0.33	0.20	0.34	0.26	0.26	0.41	-0.00	-0.27	0.77	-0.80	0.39	-0.28

Table A.5: Continued.

	TotalAeroRev13	TotalNonAeroRev13	TotalOpExp13	NetOpIncome13	Satisfaction	coordination	low	high
RegionE								
Airlines								
Routes								
Terminals								
Gates								
Runways Number								
Total pax								
Delay per flight								
Avg pax/route								
Net basic utility								
Runway usage								
Terminal usage								
Avg pax/flight								
FLT_DEP_1								
turn_delay								
Capacity (Hour)								
CUI								
Peak hour vol								
Dom pax 2014								
Int pax 2014								
Transfer pax 2014								
Public Ownership								
CHT%								
FSC%								
LCC%								
REG%								
OTH%								
TotalAeroRev13	1.00							
TotalNonAeroRev13	0.77	1.00						
TotalOpExp13	0.80	0.97	1.00					
NetOpIncome13	0.94	0.74	0.69	1.00				
Satisfaction	0.50	0.51	0.49	0.47	1.00			
Coordination	0.30	0.26	0.31	0.20	0.17	1.00		
Low	-0.26	-0.37	-0.34	-0.26	-0.12	0.02	1.00	
High	0.36	0.45	0.43	0.34	0.21	0.09	-0.89	1.00

Table A.7: Continued.

	Re- gionE	Air- lines	Routes	Ter- mi- nals	Gates	Run- ways Num- ber	Total PAX	Delay per flight	Avg pax/ route	Net Basic Util- ity	Run- way Us- age	Ter- minal Us- age	Avg pax/ flight
RegionE													
Airlines													
Routes													
Terminals													
Gates													
Runways Number													
Total PAX													
Delay per flight													
Avg pax/route													
Net Basic Utility			N/S				N/S						
Runway Usage													
Terminal Usage													
Avg pax/flight													
FLT_DEP_1										N/S			
turn.delay										N/S			
Capacity (Hour)													
CUI													
Peak Hour Vol.													
dom pax 2014													
int pax 2014													
transfer pax 2014										N/S			
Public Ownership													
CHT%			N/S		N/S								
FSC%													
LCC%		N/S											
REG%		N/S	N/S		N/S								
OTH%				N/S							N/S		
TotalAeroRev13													
TotalNonAeroRev13										N/S			
TotalOpExp13													
NetOpIncome13													
Satisfaction	N/S												
coordination	N/S												
low		N/S											
high													

Table A.8: Significance statistical tests of correlation coefficients. All correlations are significant with a 5% p-value, except where N/S is indicated.

	FLT_ DEP_1	turn_ delay	Ca- pac- ity (Hour)	CUI	Peak Hour Vol.	dom pax 2014	int pax 2014	trans- fer pax 2014	Pub- lic Own- er- ship	CHT%	FSC%	LCC%	REG%	OTH%
RegionE														
Airlines														
Routes														
Terminals														
Gates														
Runways Number														
Total PAX														
Delay per flight														
Avg pax/route														
Net Basic Utility														
Runway Usage														
Terminal Usage														
Avg pax/flight														
FLT_DEP_1														
turn_delay														
Capacity (Hour)														
CUI														
Peak Hour Vol.														
dom pax 2014														
int pax 2014														
transfer pax 2014														
Public Ownership														
CHT%						N/S								
FSC%														
LCC%	N/S	N/S												
REG%			N/S	N/S	N/S			N/S						
OTH%								N/S				N/S		
TotalAeroRev13														
TotalNonAeroRev13														
TotalOpExp13												N/S		
NetOpIncome13														N/S
Satisfaction														
coordination														
low		N/S	N/S		N/S									
high														

Table A.9: Continued.

	Total- AeroRev13	TotalNon- AeroRev13	TotalOp- Exp13	NetOpIn- come13	Satisfaction	coordina- tion	low	high
RegionE								
Airlines								
Routes								
Terminals								
Gates								
Runways Number								
Total PAX								
Delay per flight								
Avg pax/route								
Net Basic Utility								
Runway Usage								
Terminal Usage								
Avg pax/flight								
FLT_DEP_1								
turn_delay								
Capacity (Hour)								
CUI								
Peak Hour Vol.								
dom pax 2014								
int pax 2014								
transfer pax 2014								
Public Ownership								
CHT%								
FSC%								
LCC%								
REG%								
OTH%								
TotalAeroRev13								
TotalNonAeroRev13								
TotalOpExp13								
NetOpIncome13								
Satisfaction								
coordination								
low								
high								

Table A.10: Continued.

A.2 Clustering analysis

A.2.1 Methodology

There are many different ways of clustering data, each corresponding to the definition of ‘clustering’. Several methods are routinely used in the literature but the specific choice of method is always quite subjective. We decided to use a technique coming from network theory, based on modularity. If we consider a network with an adjacency matrix A and a partition P of its nodes, the modularity is defined as:

$$Q = \frac{1}{2m} \sum_{C \in P} \sum_{i,j \in C} (A_{ij} - P_{ij}), \quad (\text{A.1})$$

where P_{ij} is the expected value of the adjacency matrix for the link i, j , and m is the total weights of the links. The modularity is typically a measure of how much the nodes are tightly linked to each other within the communities (clusters), with respect to how much they are linked with the rest of the network. The null model for the matrix is usually the one proposed by Newman and Girvan [57]: $P_{ij} = k_i k_j / 2m$, which corresponds to a randomisation of the links, conserving the local strengths. There is then the need to find the partition P of nodes which maximises the modularity, and for this several algorithms exist. Here we use the Louvain method, which is very efficient and widely used [58].

It is well known that the modularity suffers from a resolution issue, but there is an easy and elegant way to circumvent this, by adding a scaling term to the null model’s matrix, i.e. $P_{ij} = \gamma k_i k_j / 2m$. If γ is high enough, typically very small communities are obtained (down to the size of one node each). A small value means on the other hand that the partition maximising the modularity is the one where all nodes are in the same partition. In between, different levels of granularity of the system are spanned. If there is again a certain degree a subjectivity in the choice of the right scale, this is strongly guided by the appearance of plateaus in the number of clusters when sweeping the scale – as shown below.

In order to use this method, the data must be organised in some type of network. A typical approach is to define a degree of similarity – or distance – between airports. Several choices are possible, but a common choice is to use the Euclidean distance on standardised data, i.e. computing:

$$d_{ij} = \sqrt{\sum_k (c_i^k - c_j^k)^2}, \quad (\text{A.2})$$

where c_i^k is the standardised value of the component k for airport i . The components

are standardised as follows:

$$c_i^k = \frac{\tilde{c}_i^k - \min_j \tilde{c}_j^k}{\max_j \tilde{c}_j^k - \min_j \tilde{c}_j^k}, \quad (\text{A.3})$$

where \tilde{c}_j^k is the initial, absolute value of the component. This means that all components span the interval $[0, 1]$. Note that different weights can also be put on different components, to better reflect either their importance or some prior knowledge of the data.

In this study we use the components of the PCA for the distances between airports, instead of the initial variables, which reflect the natural organisation of the data. Moreover, we use the relative variance weights of each component in the Euclidean distance, i.e. $d_{ij} = \sqrt{\sum_{k=0}^3 w_k (c_i^k - c_j^k)^2}$, where w_k is the ratio of variance explained with component k .

In summary, we define a certain number of components (the same that we used for the PCA), we build a network where each node is an airport and each pair of airports has links of strength $1 - d_{ij}$, we sweep the parameter γ , and for each value we compute the best partition with the Louvain algorithm. In the following subsection we show the result of the procedure.

A.2.2 Analysis of clusters

The first step is to check if there are some scales for which the system has a non-trivial number of clusters. Figure A.1 shows the existence of plateaus when sweeping the scaling parameter, more specifically a plateau with three communities and another one with four.

We then check that these plateaus actually correspond to stable partitions and not, for example, to a sequence of different partitions with the same number of clusters. We thus compute the normalised mutual information (NMI), which is a widely-used measure of how alike two partitions are. We take several values of the γ within each plateau, compute the NMI between each pair of the corresponding partitions, and average these results. The NMI score ranges between 0 and 1, where 0 indicates two totally unrelated partitions and 1 indicates that the partitions coincide exactly. In other words, if the average NMI is 1, we conclude that the partition within each plateau is unique, which is indeed the case here. Overall, this indicates ‘good’ partitions, since they are stable in a finite range of the scaling parameters. Another check of the relevance of the partition can be made by disturbing the data with some noise and observing whether the new partition is similar to the undisturbed one. The procedure applied is as follows: we first choose a level of noise (for instance $\alpha = 1\%$) and for each characteristic of each airport, we modify its value by a factor drawn randomly




Figure A.1: Number of communities in the network of airports as a function of the scaling parameter. The presence of plateaus indicates the existence of natural cluster structure at these scales.

between $1 - \alpha$ and $1 + \alpha$. We then recompute the partition (with the same value for γ) and compute the NMI between this partition and the initial one.

In Figure A.2, we show the NMI as a function of the level noise applied to the data. Clearly, for high levels of noise, the partitions are very different from the initial ones. The fact that there is a small plateau for low levels of noise is a further indication that the partition is quite robust. However, it is also clear that the clusters of the initial partition do not differ greatly, since the NMI drops quite substantially after the plateau. Overall, we can conclude that this partition is robust statistically speaking, but that the characteristics of the airports tend to be continuous rather than harshly categorical, which was as expected. Note also that a clustering analysis with many different characteristics tends to produce these types of results, because of the variance induced by each additional component.

We then inspect the partitions themselves. In Table 3.2 we display the composition of the 3-cluster partition. We firstly compare this partition with the 4-cluster partition (not shown), which is, in fact, very similar. The only difference is the presence of the new cluster containing two airports, Copenhagen and Vienna, otherwise in cluster number 2. Since the operational meaning of this small cluster is not obvious, we focus in this study on the 3-cluster partition.




Figure A.2: Average Normalised Mutual Information (NMI) between partitions based on disturbed data and the original 3-clusters partition, as a function of the level of noise applied to the data. The error bars represent the standard deviations over the 100 realisations of noise.

B The model

B.1 Calibration

B.1.1 The problem of delay

One of the core functions of the model is the relationship between delay and traffic. It is clear that higher traffic lead to higher delays on average, however the link is not so easy to infer. If at first we used some real delay distribution taken from CODA data, we realised that we needed the full distribution of delays during the day¹, which is only available to us through DDR data. Unfortunately, DDR data only includes M1 and M3 departures times, which means that we have only access to some delay after regulation and updates from the airlines. In order to see if both delays were closely related or not, we produced the plot presented in Figure B.1 (here for LFPG). Fortunately, both distributions are not so far from each other. In fact, the main differences are in the negative delay, essentially because a high negative delay from updated flight plans turns into a null delay for CODA, which is fairly normal. The second difference is the average value, which is around 9.5 minutes per flights in DDR data and around 3.5 minutes in CODA data. We could have shifted the former to match the latter, but decided that it would be lengthy to do for all airports, and is not really meaningful. As a consequence, we decided to use the measure of delay from DDR as it is in the following.

B.1.2 Delay and capacity

Once we had the above approximated measure of delay, we tried to link the delay at an airport with the capacity and the traffic. In order to do this, we performed a regression for each airport of the delay against the traffic, as described in Section 3.3. Figure B.2 shows the evolution of the delay as a function of traffic for LFPG, and the corresponding exponential fit. The fitted function has the form $\delta = 120(\exp(N_f/C) - cc)$, with cc and C being the two regressed parameters. C can be interpreted as the

¹As highlighted in the main text and in the following, we need the full distribution because of the non-linearity of the cost of delay for airlines.




Figure B.1: Distribution of delays in CODA data ('real' delays) and DDR data for LFPG. The data has been truncated in the $[-15, 120]$ region (following truncation from CODA data).




Figure B.2: Delay in seconds as a function of the traffic, for each hour of the day. The solid line is an exponential fit.

capacity, since it is the typical traffic for which the delay significantly differs from its original value. cc is the delay when the traffic tends to zero, which is slightly negative for most of the airports. Note that a linear fit is sometimes as good as an exponential one, as shown below. Also the delays are still quite small (less than 10 minutes), since they are averaged over an hour. We will tackle the issue of very high delays for airlines in the discussion that follows. Finally, note that the exact resolution for the fit is not obvious *a priori*, as our definition of capacity is now detached of any time period other than the averaging window. To tackle this problem, we performed for each airport a maximisation of the fit goodness by sweeping the resolution. Interestingly, the optimal resolution is not the same for every airport, which may reflect different operational choices.

Table B.1 shows the results of the fits for all the airports. For most of them, the fit is good, except for a couple (LIRF and EIDW). In order to use the capacity in the model, we rescale it to have the equivalent hourly capacity (\mathcal{C}_{hour} in the table). Moreover, we decided to use the clusters found in the data analysis section in order to get an average capacity. Indeed, it is quite clear that the fitted capacities fluctuate significantly within one cluster, so we are more confident in using an average one, since the airports in the same cluster should have similar operations. The last column of the table show these averages, which are used in the model.

B.1.3 Cost of delay

The cost of delay is quite complex to assess, given the non-linearity of the cost as a function of delay, and of the large distribution of delay. The basic cost per flight is shown in equation 3.1. However, the average delay in this equation cannot simply be used in order to have the average cost. The easiest way to explain this is to acknowledge that, even when the average delay is null, the average cost is not, because some flights have a positive delay (i.e. arrive early) whereas others have a negative delay and do not increase profits. As a result, a full distribution of delay was considered.

In order to do this, and maintain a relatively simple model, we used a lognormal fit of the distribution of delay. Since for each hour of the day, the distribution is likely to differ, we fitted the distribution of delay of each hour independently. Note that we used the whole month of data by aggregating each hour of each of the month. Figure B.3 shows an example of a fit for 9.00 am for LFPG. The fit is far from perfect, but roughly matches the data, and thus allows us to have a computation of the expected. Once we have fitted the probability of delay $p(\delta)$, the expected cost of delay \bar{c}_d is:

$$\bar{c}_d = \int c_d(\delta t) p(\delta t) d(\delta t)$$

By performing this procedure for each hour of the day, the total cost of delay for

	Res.	\mathcal{C}_{res}	cc	R^2	Cluster	\mathcal{C}_{hour}	\mathcal{C}_{com}
EBBR	74	28.296911	-1.429895	0.709285	2	22.943441	29.393561
EDDL	20	11.831204	-4.072900	0.617365	2	35.493611	29.393561
EGCC	81	20.052234	0.451433	0.642084	2	14.853507	29.393561
EIDW	94	71.311668	-3.025547	0.056397	2	45.518086	29.393561
EKCH	18	13.500190	-3.626090	0.870522	2	45.000634	29.393561
ENGM	46	26.990120	-0.837604	0.382568	2	35.204504	29.393561
ESSA	50	25.630759	-0.869383	0.564417	2	30.756911	29.393561
LOWW	29	16.600777	-1.437840	0.881326	2	34.346435	29.393561
LPPT	57	15.721970	-3.438412	0.732810	2	16.549442	29.393561
EDDF	35	21.956503	-2.452524	0.719974	1	37.639719	35.587806
EDDM	24	29.398295	-3.367913	0.599480	1	73.495738	35.587806
EGKK	46	15.507388	-1.136089	0.780133	1	20.227027	35.587806
EGLL	29	13.882645	-0.810045	0.821380	1	28.722714	35.587806
EHAM	46	38.471294	-2.787346	0.494721	1	50.179949	35.587806
LEMD	46	22.940405	-2.734194	0.521661	1	29.922268	35.587806
LFPG	47	32.313652	-1.483228	0.628944	1	41.251471	35.587806
LIRF	91	438.107591	-3.775606	0.000976	1	288.862148	35.587806
LSZH	70	30.412924	-2.072356	0.640382	1	26.068220	35.587806
LTBA	84	17.896407	0.334670	0.817061	1	12.783148	35.587806
EDDH	91	21.696123	-1.831217	0.522407	0	14.305136	17.813555
EDDK	22	14.923674	-3.962501	0.327822	0	40.700929	17.813555
EFHK	35	17.240393	-1.771224	0.660697	0	29.554959	17.813555
EGBB	53	9.316811	-0.074231	0.886582	0	10.547333	17.813555
EGSS	67	19.592047	-0.633787	0.544832	0	17.545117	17.813555
EPWA	80	17.771291	-2.708330	0.656717	0	13.328468	17.813555
LFMN	91	13.716757	-2.999108	0.840659	0	9.044016	17.813555
LGAV	74	16.173872	-0.549414	0.902492	0	13.113950	17.813555
LHBP	45	17.058659	-2.998551	0.496665	0	22.744879	17.813555
LKPR	63	21.392047	-2.366516	0.685793	0	20.373378	17.813555
LPPR	82	6.410945	-0.009339	0.322322	0	4.690935	17.813555

Table B.1: Results of the exponential fits of the delay against the traffic. The first column is the resolution (in minutes) which gives the best fit. The second and third columns are the corresponding values of the fitting parameters, and the fourth one is the corresponding coefficient of determination. The fifth column is the cluster ownership (as found in the data analysis part). The sixth column is the capacity recomputed for one hour, which is the resolution of the model. The last column features the average capacity in each cluster, which is taken as the input by the model.




Figure B.3: Histogram of the delays for LFPG at 9am, with a lognormal fit.

the airlines in one day is obtained by summing all the individual contributions. This ‘corrected’ cost is indeed significantly different from the cost of the average delay, as shown on Figure B.4. Note that the correction itself depends on each airport, because the exact distribution of delay also does. This is the curve we directly use in the model.

B.2 Model implementation

In this section, we present the specific equations that we used in the model. The first two equations we used are related to the delay-capacity relationship and the cost of delay, as presented previously. The other equations are detailed below.

B.2.1 Probability of operation

The core principle of the model is that the airline has a probability P_A of operating a flight based on the expected cost of delays, which reads:

$$P_A = 2 \frac{1}{1 + e^{c_d/s}},$$

where c_d is the expected cost of delay computed along the lines presented previously, and s is the smoothness parameter, discussed in the main text. This choice is motivated by the fact that the probability is linked to some form of utility function for




Figure B.4: Correction of cost of delay for airline. The blue points are the cost of the average delay, the violet ones are the expected costs taking into account the distribution of delay, and the solid red line is a quadratic fit of the latter.

the airline, taking into account other (strategic) parameters (as described above). It allows us to have a smooth function which varies continuously between 0 and 1, and to have a risk aversion of the agent which can directly be linked to the parameter s – henceforth referred to as the ‘smoothness’ of the decision. Indeed, when s is sufficiently small, the airline takes harsh decisions, switching from operating to non-operating the route when the costs are driven high enough. Note that, in fact, we should strictly be referring to net revenue contribution to the network, since airlines will tolerate loss-making legs that have a net benefit to the system.

Moreover, we are able to introduce with this function an element of prospect theory, in the sense that the utility of the airline does not depend only on its revenues, but also on a pre-determined level r_0 (some kind of ‘anchoring’), by using $c_d + r_0$ instead of simply c_d . This parameter could include the direct revenues from the passengers (i.e prices of the tickets) and other costs linked to the operations. In the final version of the model we decided not to use this parameter because it was hard to calibrate and did not bring interesting features (we would need to consider a full distribution of r_0 for this, which is beyond the scope of this research). Finally, this function mimics standard functions from prospect theory, since it is convex in the positive region (with revenues greater than the value anchor) and concave otherwise.

B.2.2 Implicit equation of delay

One important feature of the model is that it is auto-consistent for the delays, i.e. the delays in the output are exactly the right level to match the actual traffic, which in turn sets the average delay through the delay-capacity relationship. In other words, setting a distribution of delays fixes the actual traffic through the use of probability P_A , which in turn fixes the delay at each hour of the day through the capacity-delay relationship. In order to solve this circular issue, we need to solve the implicit equation:

$$\delta = 120 (\exp(P_A(\delta t)T/C) - cc) \quad (\text{B.1})$$

$$= 120 \left(\exp \left(2 \frac{1}{1 + e^{c_d(\delta t)/s}} \frac{T}{C} \right) - cc \right), \quad (\text{B.2})$$

$$(\text{B.3})$$

using the corrected function c_d of the delay, as described above. T is the demand traffic for this specific hour, i.e. the traffic extracted from data multiplied by β , as explained in the main text. This equation is solved numerically using the Brent method. The convergence is assured by the fact that the right hand side is monotonically decreasing with delay, and is strictly superior to 0 when $\delta t = 0$. So, there is always a solution, and only one, which is why we use a local optimisation algorithm.

B.2.3 The airport

Regarding the airport, we assume that its revenues come from aeronautical revenues and non-aeronautical revenues. The former depends on the airport charges P and the potential number of flights operated N , with:

$$r_{A,\text{aero}} = PNP_A.$$

The latter is directly linked to the number of passengers:

$$r_{A,\text{non-aero}} = l_f w N P_A,$$

where l_f is the average load factor and w is the average revenue coming from each individual passenger. In the final version of the model, the latter is a constant too.

Finally, we consider the expenses of the airport. Since we are interested in capacity-related costs, we choose a very simple form for the cost, linear with the extra capacity. The cost function reads:

$$c_{inf} = \alpha(\mathcal{C} - \mathcal{C}_{init}) + c_{init}, \quad (\text{B.4})$$

where $\mathcal{C} - \mathcal{C}_{init}$ represents the increase in capacity desired by the airport, α is the marginal cost per unit of capacity, and c_{init} is the initial total expense of the airport – i.e. the expenses found in the data.

In order to compare the consequences of the model for the different actors, we also compute a utility for the passengers:

$$u_p = -v\delta t, \quad (\text{B.5})$$

where v is the average value of time of the passengers. This is a very crude approximation, and the computed utility has no intrinsic meaning, but is nevertheless useful to compare different situations where delays are different and passengers are likely to be satisfied at different levels. Note that we do not compare the results of the model for different airports and thus the satisfaction coming from the airport does not need to be estimated.

B.3 Exploratory analysis

This section shows some details of the results presented in 4.3.

The equation we used to modify the non-aeronautical revenues of the airport is the following:

$$w(\delta t) = w_{init} + w_{shop}(\delta t) + w_{sat}(\delta t).$$

It is composed of two components, plus the constant part w_{init} , coming directly from the calibration. The first one is increasing with δt , and represents the tendency of people to spend more when they have more time ('more shopping time'). The second component is decreasing with δt and represents the tendency of people to spend more when their satisfaction is higher ('better shopping time'). Since we assumed that the second component has an effect only when the delay is very high, whereas the first one is present as soon as waiting increases, we use a quadratic and linear form for these. The first reads:

$$w_{shop}(\delta t) = t_e \frac{\delta t - \delta t_{init}}{120} w_{init},$$

and the second:

$$w_{sat}(\delta t) = \begin{cases} s_e \left(\frac{\delta t - \delta t_{init}}{120} \right)^2 w_{init} & \text{if } \delta t < \delta t_{init} \\ -s_e \left(\frac{\delta t - \delta t_{init}}{120} \right)^2 w_{init} & \text{otherwise.} \end{cases}$$

In these equations, δt_{init} represents the initial delay (the delay just after calibration). t_e and s_e are the weights given to each component, and are considered as free parameters. Finally, we cap the value w , to avoid infinite revenues. The corresponding variable is noted *cap* in the main text and is set to 24 euros to produce the results (to be compared to the 17.7 euros after calibration).

As a result of this formulation, we are able to produce a non-trivial curve, shown in Figure B.5, which in turns contribute to the net income curve shown in the main text.




Figure B.5: Evolution of the function w with the delay, with $\delta t_{init} = 9.5$, $w_{init} = 17.8$, $t_e = 15$, and $s_e = 1000$.

C Sensitivity analysis

In this annex, we briefly show the results of a sensitivity analysis performed on a calibrated example of a large, European hub airport.

Since we have only one free parameter left in the model, we simply sweep it and see how the calibrated parameters changes. In Figure C.1, we show the evolution of the average delay in the output and the revenues of the airlines (in fact, only the cost of delay, negatively counted). These two outputs are the ones which are of interest, all others being fixed (e.g. the revenues per passengers) or trivially related to them (e.g. the utility of the passenger). Both quantities are changing with the smoothness,




Figure C.1: Evolution of the average delay (left) and revenues of airlines (right) in the calibrated model for various values of the smoothness parameter s .

but not drastically. For instance the delay goes from around 9 minutes per flight up to 11.7 minutes, which is a fairly narrow window, even though it is not negligible. It is worth noting that the actual value of the delay for the calibrated airport is 9.5 minutes in the data, which means in fact that we could calibrate this last parameter to fit the average delay. We did not do it for technical reasons, but in the main text we chose $s = 500$, which gives a delay close to 9.5. We are thus confident that the results presented in the main text are sufficiently reliable with regard to the parameters.

Bibliography

- [1] P. McCarthy, US airport costs and production technology a translog cost function analysis, *Journal of Transport Economics and Policy* 48 (3) (2014) 427–447.
- [2] S. Gleave, Prepared for the UK CAA (2012).
- [3] A. Cook, G. Tanner, European airline delay cost reference values, updated and extended values, version 4.1, <https://www.eurocontrol.int/publications/european-airline-delay-cost-reference-values> (2015).
- [4] CAA, Licence granted to heathrow airport limited by the civil aviation authority under section 15 of the civil aviation act 2012 on 13 february 2014 (2016).
- [5] D. G. D. l’Aviation Civile, Economic regulation agreement between the government and Aeroports de Paris 2016-2020 (2015).
- [6] T. D’Alfonso, C. Jiang, Y. Wan, Airport pricing, concession revenues and passenger types, *Journal of Transport Economics and Policy* (1) (2013) 71–89.
- [7] E. Torres, J. S. Domínguez, L. Valdes, R. Aza, Passenger waiting time in an airport and expenditure carried out in the commercial area, *Journal of Air Transport Management* 11 (6) (2005) 363–367.
- [8] M. Geuens, D. Vantomme, M. Brengman, Developing a typology of airport shoppers, *Tourism Management* 25 (5) (2004) 615–622.
- [9] J. Castillo-Manzano, Determinants of commercial revenues at airports: lessons learned from Spanish regional airports, *Tourism Management* 31 (2010) 788–796.
- [10] Y. Lin, C. Chen, Passengers’ shopping motivations and commercial activities at airports – the moderating effects of time pressure and impulse buying tendency, *Tourism Management* 36 (2013) 426–434.
- [11] F. Fuerst, S. Gross, U. Klose, The sky is the limit? the determinants and constraints of European airports’ commercial revenues, *Journal of Air Transport Management* 17 (5) (2011) 278–283.

- [12] V. Fasone, L. Kofler, R. Scuderi, Business performance of airports: Non-aviation revenues and their determinants (2016).
- [13] S. Appold, J. Kasarda, The appropriate scale of us airport retail activities (2006).
- [14] Y. Wan, C. Jiang, A. Zhang, Airport congestion pricing and terminal investment: Effects of terminal congestion, passenger types, and concessions, *Transportation Research Part B: Methodological* 82 (2015) 91–113.
- [15] Airports Council International, Does passenger satisfaction increase airport non-aeronautical revenue? A comprehensive assessment, research report (2016).
- [16] CAA, Heathrow: Market power assessment (2012).
- [17] R. Merkert, G. Assaf, Using DEA models to jointly estimate service quality perception and profitability – evidence from international airports, *Transportation Research Part A: Policy and Practice* 75 (2015) 42–50.
- [18] Y. Xiao, X. Fu, A. Zhang, Demand uncertainty and airport capacity choice, *Transportation Research Part B: Methodological* 57 (2013) 91–104.
- [19] C. Barnhart, D. Fearing, A. Odoni, V. Vaze, Demand and capacity management in air transportation, *EURO Journal on Transportation and Logistics* 1 (1-2) (2012) 135–155.
- [20] J. Martín, A. Voltes-Dorta, The dilemma between capacity expansions and multi-airport systems: Empirical evidence from the industry’s cost function, *Transportation Research Part E: Logistics and Transportation Review* 47 (3) (2011) 382–389.
- [21] J. Daniel, Benefit-cost analysis of airport infrastructure: the case of taxiways, *Journal of Air Transport Management* 8 (3) (2002) 149–164.
- [22] J.-D. Jorge, G. de Rus, Cost-benefit analysis of investments in airport infrastructure: a practical approach, *Journal of Air Transport Management* 10 (5) (2004) 311–326.
- [23] M. Gelhausen, P. Berster, D. Wilken, Do airport capacity constraints have a serious impact on the future development of air traffic?, *Journal of Air Transport Management* 28 (2013) 3–13.
- [24] P. Berster, M. Gelhausen, D. Wilken, Is increasing seat capacity common practice of airlines at congested airports?, *Journal of Air Transport Management* 46 (2013) 1–25.

- [25] A. Zhang, A. Czerny, Airports and airlines economics and policy: an interpretive review of recent research, *Economics of Transportation* 1 (1) (2012) 15–34.
- [26] IATA, Peak/off-peak charges.
- [27] A. Czerny, A. Zhang, Airport congestion pricing and passenger types, *Transportation Research Part B: Methodological* 45 (3) (2011) 595–604.
- [28] N. Adler, P. Forsyth, J. Mueller, H.-M. Niemeier, An economic assessment of airport incentive regulation, *Transport Policy* 41 (2015) 5–15.
- [29] M. Madas, K. Zografos, Airport capacity vs. demand: Mismatch or mismanagement?, *Transportation Research Part A: Policy and Practice* 42 (1) (2008) 203–226.
- [30] E. Verhoef, Congestion pricing, slot sales and slot trading in aviation, *Transportation Research Part B: Methodological* 44 (3) (2010) 320–329.
- [31] B. Desart, D. Gillingwater, M. Janic, Capacity dynamics and the formulation of the airport capacity/stability paradox: a European perspective, *Journal of Air Transport Management* 16 (2) (2010) 81–85.
- [32] A. Cook, G. Tanner, S. Anderson, Evaluating the true cost to airlines of one minute of airborne or ground delay, University of Westminster, for EUROCONTROL Performance Review Commission, Brussels, (2004).
- [33] E. Gilbo, Airport capacity: Representation, estimation, optimization, *IEEE Transactions on Control Systems Technology* (3) (1993) 144–154.
- [34] FAA, Economic values for faa investment and regulatory decisions, a guide (2014).
- [35] A. Zhang, Y. Zhang, Airport capacity and congestion pricing with both aeronautical and commercial operations, *Transportation Research Part B: Methodological* 44 (3) (2010) 404–413.
- [36] International Transport Forum, Expanding airport capacity under constraints in large urban areas, discussion paper, (2013), p. 242.
- [37] A. Montreal, Airport economics survey (2015).
- [38] Z. Lei, A. Papatheodorou, E. Szivas, The effect of low-cost carriers on regional airports’ revenue: evidence from the UK in Forsyth P., Gillen D., Muller J. and Niemeier H.-M. (eds), *Airport Competition: The European Experience*, Aldershot: Ashgate (2010).

- [39] V. Liebert, H. Niemeier, A survey of empirical research on the productivity and efficiency measurement of airports, *J. Transp. Econ and Policy* 47 (2) (2013) 157–189.
- [40] E. Pels, P. Nijkamp, P. Rietveld, Inefficiencies and scale economies of European airport operations, *Transportation Research Part E: Logistics and Transportation Review* 39 (5) (2003) 341–361.
- [41] J. Martin, A. Voltes-Dorta, International airports: economies of scale and marginal costs 47 (1) (2010) 5–22.
- [42] EC, Communication from the commission: Guidelines on state aid to airports and airlines, *Official Journal C99*, 4 April (2014).
- [43] CAA, Economic regulation at new runway capacity, cap 1279, london, caa (2015).
- [44] K. Van Dender, *J. Urban Econ.* 62 (2007) 317–336.
- [45] Y. Choo, Factors affecting aeronautical charges at major us airports, *Transportation Research Part A* 62 (2014) 54–62.
- [46] G. Bel, X. Fageda, Privatization, regulation and airport pricing: an empirical analysis for Europe, *Journal of Regulatory Economics* 37 (2) (2010) 142–161.
- [47] V. Bilotkach, J. Clougherty, J. Mueller, Z. A, Regulation, privatization, and airport charges: panel data evidence from European airports, *Journal of Regulatory Economics* 42 (1) (2012) 73–94.
- [48] X. Fu, M. Lijesen, T. H. Oum, An analysis of airport pricing and regulation in the presence of competition between full service airlines and low cost carriers, *Journal of Transport Economics and Policy (JTEP)* 40 (3) (2006) 425–447.
- [49] CAA, Stansted: Market power assessment - annex 3 (2013).
- [50] SEO, Impacts of expanding airport capacity on competition and connectivity, itf (2014).
- [51] D. Starkie, G. Yarrow, Why airports can face price-elastic demands: Margins, lumpiness and leveraged passenger losses, *International Transport Forum discussion paper* 23 (2013).
- [52] EUROCONTROL, Coda digest: All-causes delay and cancellations to air transport in Europe (2014).

- [53] European Commission, 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, 17 february 2004, 1-7 (2004).
- [54] J.-O. Kim, C. Mueller, Factor Analysis – Statistical Methods and Practical Issues, SAGE Publications, Newbury Park, California, (1978).
- [55] EUROCONTROL, Standard Inputs for EUROCONTROL Cost Benefit Analysis (Edition Number: 7.0, November 2015), accessed on 28.06.2016, <https://www.eurocontrol.int/sites/default/files/publication/files/standard-input-for-eurocontrol-cost-benefit-analyses-2015.pdf> (2015).
- [56] <http://www.mathworks.com/products/compiler/mcr>.
- [57] M. Newman, M. Girvan, Finding and evaluating community structure in networks, Phys Rev E 69 (2003) 026113.
- [58] V. Blondel, J. Guillaume, R. Lambiotte, E. Lefebvre, Fast unfolding of communities in large networks, JSTAT (2008) P10008.