Jan-Rasmus Künnen

Advanced demand-capacity balancing mechanisms
to improve performance of European air traffic networks

Dissertation
for obtaining the degree of Doctor of Business and Economics
(Doctor rerum politicarum - Dr. rer. pol.)

at WHU – Otto Beisheim School of Management

May 12, 2023

First Advisor: Prof. Dr. Arne K. Strauss

Second Advisor: Prof. Dr. Stefan Spinler

It is not because things are difficult that we do not dare, but because we do not dare that things are difficult.

– *Seneca the Younger*

Acknowledgements

Working on this dissertation has been one of the most challenging and rewarding experiences in my life. It has helped me grow both professionally and personally. Throughout my time at WHU, I had the privilege to be accompanied by amazing people. I would therefore like to acknowledge a few exceptional individuals without whose help these pages would have remained unwritten.

First, I am extremely grateful to my first advisor *Prof. Dr. Arne K. Strauss*, who has always given me the confidence and support to take the necessary steps in my dissertation. It is through your professional guidance and dedication that every hard problem became computationally tractable, and that enabled me to venture into the realms of mathematical optimization. I would also like to thank my second advisor *Prof. Dr. Stefan Spinler*. Your practical advise has helped me shape and improve this dissertation. My endeavor into European air traffic management would not have been as fruitful (and pleasant) without the help of the CADENZA research team, most notably *Nikola, Radosav* and *Frank*. Your experience proved invaluable as a sounding board for new ideas, and the collaboration with each and every one of you lightened up long modeling days.

The dissertation process may come to an end, but friends are there to stay. My special thanks go to current and former colleagues and friends: *Sebastian*, for pulling me to the WHU in the first place; *Johannes*, for nudging me to take flying into practice; *Simon*, for keeping me return to Hamburg frequently; and *Jasper, Patrick* and *Thomas* for continuously pushing my (physical) boundaries across the Alps and Dolomites. I am looking forward to spending more time with all of you in the future.

Most importantly, I would like to thank my parents for unconditionally supporting all my life decisions, no matter how near of far from home they took me. Finally, I would like to thank *Annika*, my significant other and partner. Our time together may have started during this endeavor, but it will last for many more!

Vallendar, August 2022

Jan-Rasmus Künnen

Contents

List of Figures

List of Tables

List of Abbreviations

ACC	Area control center
ANSP	Air navigation service provider
ATC	Air traffic control
ATCO	Air traffic controller
ATFM	Air traffic flow management
ATM	Air traffic management
AU	Airspace user
DCB	Demand-capacity balancing
MCKP	Multi-choice knapsack problem
MMKP	Multi-dimensional multi-choice knapsack problem
NM	Network manager

Notation

Table 1. Overview of notation across dissertation.

Sets:	
$f \in \mathcal{F}$	Finite collection of flights
$r \in R^f$	Finite set of re-routing and delay options available to flight $f \in \mathcal{F}$
$u \in U$	Set of time periods covering the day of operation
$a \in A$	Set of airspaces
$c \in C^a$	Set of configurations for airspace a
$l \in L^c$	Set of collapsed sectors corresponding to configuration c
$e \in E^l$	Subset of elementary sectors forming sector l
$S \in \mathcal{S}$	Finite set of scenarios for demand and capacity uncertainty
Parameters:	
$\kappa = (\kappa_l)$	Nominal sector capacity for each elementary or collapsed sector l
$\kappa^S = (\kappa_l^S)$	Maximum capacity of sector l under scenario S
$\bar{h} = (\bar{h}_{ac})$	Sector-hours consumed by airspace a in configuration c
$k = (k_{acu})$	Capacity shortage in configuration c of airspace a at time unit u
$d = (d_r^f)$	Displacement cost of route r for flight f
$b = (b_{freu})$	Indicates whether flight f on route r uses sector e at time u

Table 2. Overview of notation for Chapter 2.

Sets:	
$t \in T$	Set of time periods covering the booking horizon
$z \in Z$	Product types among which AUs can choose
$f \in F_t$	Set of flights that have arrived to booking process until t
$f \in F^G$	Set of regular (or scheduled) flights
$f \in F^N$	Set of non-scheduled flights
W^S	Materialized demand and capacity under scenario S
$p \in Pr$	Finite set of relative price points for trajectory products
Parameters:	
$H = (H_a)$	Capacity budget for airspace a, given from strategic phase
Decision variables:	
$\mathbf{x} = (x_r^f)$	Indicates whether flight f is assigned to route r
$\mathbf{y} = (y_{acu})$	Indicates whether configuration c is open in airspace a at time u

Table 3. Overview of notation for Chapter 3 and 4.

Sets:	
F^S	Materialized demand under scenario S (i.e., flights)
W^S	Materialized capacity under scenario S
Parameters:	
$\gamma = (\gamma_a)$	Unit cost of one sector hour for airspace a
Decision variables:	
$\mathbf{x} = (x_a)$	Capacity level (in sector-hours) for airspace a
$\mathbf{x^0} = (x_a^0)$	Virtual capacity (in sector-hours) provided by airspace $a \in A$
$\mathbf{h^0} = (h_a^0)$	Virtual capacity (in sector-hours) deployed in airspace $a \in A$
$\mathbf{y} = (y_r^f)$	Indicates whether flight f is assigned to route r
$\mathbf{z} = (z_{acu})$	Indicates whether configuration c is open in airspace a at time u

Chapter 1

Introduction

The field of air traffic management (ATM) is concerned with ensuring a safe passage of flights across the airspace. Operationally, the job is conducted by air traffic controllers (ATCOs) that monitor flight paths and steer flights through the network. Since ATCOs are expensive resources which require years of training before becoming operational, it is an important task of ATM to balance the demand for air traffic with the capacity of the network (in terms of ATCOs) at all times. This is especially relevant in European ATM, where a highly fragmented airspace and large traffic volumes exacerbate the problem. In the years leading up to the Covid pandemic, the European airspace showed particularly high demand-capacity imbalances. In a typical summer week in 2017, demand for air traffic services exceeded capacity 7% of the time, leading to ATM-related delays costing around €550 million that year (Eurocontrol, 2018). At the same time, the load of airspace sectors was reported to be below 60% of capacity half of the time (Eurocontrol, 2018), leading to large opportunity cost from spare capacities. These imbalances of capacity and demand occur because capacity decisions remain largely uncoordinated, and because once capacity levels are set, there are limited ways to adjust them to demand (which is volatile and only materializes over time).

In order to handle all traffic in the case of insufficient network capacity, ATCOs may impose ground delays on flights that would otherwise cross congested airspace sectors. This often leads to unintended consequences: To avoid the delays imposed on them, some airspace users (AUs) may instead choose to fly longer routes that avoid the congested sectors, which increases not only the flying time, but also fuel burn and flight emissions. According to Eurocontrol, 8.6%-11.2% of aviation emissions in Europe can be attributed to the fuel-inefficient routing of flights in 2019 (Eurocontrol, 2020c). The European Green Deal requires that transport emissions be reduced by 90% until 2050 compared to 1990 levels (European Com-

mission, 2021). Achieving this target will require improvements on all fronts: next to the use of more fuel-efficient aircraft and the development of sustainable aviation fuels (to substitute kerosene), improvements in ATM are expected to contribute a 6% reduction in annual aviation emissions (Destination2050, 2021).

This dissertation is concerned with developing demand-capacity balancing mechanisms for European ATM that improve network performance and reduce flight emissions. Rethinking the existing ATM process requires concepts for stronger collaboration and coordination across stakeholders in Europe, as well as mathematical models that help make the complex decision involved in the process.

1.1 Research questions and contributions

The main challenges involved in demand-capacity balancing can be best understood along the three phases of the ATM value chain: strategic, pre-tactical and tactical phase. In the strategic phase, which takes place up to 6 months before the day of operation, the capacity for each airspace needs to be decided. The capacity of an airspace is given in terms of the number of ATCOs employed on the day of operation and ultimately influences how many flights can be routed through the airspace that day. In the current process, each air navigation service provider (ANSP) decides autonomously on the capacity level in each of the airspaces that it controls, with limited guidance by the Network Manager (NM). In the pre-tactical phase, the airspace users (AUs) submit their flight intentions to the NM, indicating when and where in the network they want to fly. Currently, these intentions are submitted rather late in the process and often occur only several hours before departure. In the tactical phase, the so-called sector opening scheme and the routing of flights need to be decided - based on the strategic capacity levels and the submitted flight intentions. The sector opening scheme governs which airspace sectors should be opened at each time, which depends on the number of ATCOs employed in the airspace. While the sector-opening scheme of each airspace is again determined autonomously by the ANSP responsible for the airspace, the routing of each flight is decided jointly by the NM and the AU operating the flight. The objective of the process is to determine airspace capacities and flight routings such that total capacity and displacement costs (i.e., costs incurred from delaying or re-routing of flights) are minimized.

The outlined demand-capacity balancing process presents three key challenges. First, the capacity levels in the strategic phase need to be determined while demand (particularly from non-scheduled flights) is unknown and sector capacities may still be impacted by unexpected

shortages of ATCOs (e.g., due to sickness) or adverse weather on the day of operations. Second, the flight intentions submitted by AUs in the pre-tactical phase may not be realized with the available network capacities so that flights need to be re-routed or delayed. Third, flight emissions are not explicitly considered in the capacity or demand decisions, which may lead to inefficient routings of flights resulting either from capacity shortages or deliberate route choices.

These challenges are addressed in the dissertation and summarized in the following three research questions:

1. How do we determine airspace capacities in the strategic phase, given uncertainty in traffic and capacity provision, such that total capacity and expected displacement costs are minimized?

2. How do we steer demand in the pre-tactical phase, given the strategic capacity levels, such that actual displacement costs (from delays and detours) are minimized?

3. How can the proposed demand-capacity balancing mechanisms help save flight emissions across the ATM value chain?

To tackle the research questions, I first develop a routing heuristic based on the Multi-dimensional multi-choice knapsack problem (MMKP) to solve the integrated sector-opening and routing problem in the tactical phase of ATM, which I show is $\mathcal{NP}$-hard. Based on this heuristic I develop an efficient approach to determine the opportunity cost generated by offering a certain trajectory product for each flight, which is used as the dynamic price for this product. As a key contribution to existing research, the procedure runs sufficiently fast to determine dynamic prices in real-time. This allows us to evaluate the potential of dynamically-priced flexible trajectory products in steering air traffic demand in pre-tactical ATM – a novelty in ATM research. Furthermore, the mechanism allows us to compare the performance of flexible trajectory products with two more extreme settings where either the NM or the AUs decide autonomously on the routing of flights through the network.

The routing heuristic also builds the foundation for improved capacity planning in strategic ATM. For this purpose, the heuristic is used to generate displacement cost observations for various traffic and capacity scenarios. These cost observation then feed into a regression that approximates the relation between capacity shortage of each airspace and expected displacement cost in the network. The regression is used to estimate displacement costs in each iteration of an exploration-exploitation algorithm, which results in a capacity decision with lowest network cost across a range of capacity and traffic scenarios. In contrast to existing

capacity models, this allows us to determine stochastically-optimal capacity levels, given the uncertainty in demand and capacity provision. The capacity planning model is also adjusted to determine the 'optimal' level of capacity to be provided virtually (for capacity sharing) and thus the number of ATCOs to be trained for working also outside their local airspace. As a main contribution, this mechanism allows us to assess the viability and effectiveness of cross-border capacity sharing, which existing capacity models are not capable of.

Finally, both the routing heuristic and the capacity planning model are used to assess the potential of two demand-capacity balancing mechanisms to reduce aviation emissions in Europe. The flight emissions generated by a certain routing of flights are estimated based on the fuel efficiency of the aircraft operating the flight and a well-established parameter governing the relationship between fuel and emissions.

Overall, the dissertation adds to existing research by addressing a timely and relevant problem setting with novel solution approaches. The travel situation at airports and in the European airspace in summer 2022, with extraordinary levels of flight cancellations and delays, has shown how sensitive air traffic operations have become to unexpected changes in demand. It has also highlighted the importance of effective demand and capacity management mechanisms to ensure a more stable and resilient network performance in the future. To resolve the problem, I develop in this dissertation a set of new and practicable mechanisms that balance capacity with demand for air traffic services using optimization and simulation optimization techniques.

1.2 Structure of the dissertation

The dissertation is structured along three papers, each of which addresses one of the above-mentioned research questions. In the first paper I propose a dynamic pricing mechanism that aims at steering demand for ATM services during the pre-tactical phase such that expected network performance is improved. The second paper deals with the capacity planning problem in the strategic phase of ATM, for which capacity sharing among airspaces is proposed and its performance assessed in a realistically-sized simulation study. In the third paper, I propose a network-oriented approach to capacity planning (in the strategic phase) and trajectory-independent ATM service charges (in the pre-tactical phase) to reduce aviation emissions in Europe.

Chapter 2 presents the first paper on the topic of demand management in the pre-tactical phase of ATM. It addresses research question 2 by extending dynamic pricing to the field of

	Objective	Methodology	Case study data	Contributions
Chapter 2: The value of flexible flight-to-route assignments in pre-tactical air traffic management	Develop dynamic pricing mechanism for flexible trajectory products that helps steer demand for ATM services	▪ Dynamic programming ▪ Optimization heuristics ▪ Insertion heuristic	Artificial network with 200 flights across 5 airspaces and 1 hour interval	I. Formulating Dynamic Trajectory Pricing (DTP) problem as dynamic program with hard boundary condition, providing bounds for value function approximation II. Routing heuristic that delivers good displacement cost estimates while substantially reducing solution time III. Dynamic pricing mechanism for flexible trajectory products that delivers significantly better network performance than current demand management
Chapter 3: Cross-border capacity planning in air traffic management under uncertainty	Develop capacity planning mechanism that deals with uncertainty and can be extended to model cross-border capacity sharing	▪ Simulation optimization ▪ Optimization heuristics ▪ Least-squares regression	Real network data from Eurocontrol with 4,000 flights across 15 airspaces and 6 hour interval	I. Formulating strategic capacity planning problem with capacity sharing as two-stage newsvendor problem II. Simulation optimization approach that determines capacity decisions based on *stochastic* performance, delivering significantly lower network cost than deterministic benchmark III. First to assess value of capacity sharing in large case study on real network with different design options
Chapter 4: Leveraging demand-capacity balancing to reduce flight emissions and improve network performance	Evaluate potential of demand and capacity manage-ment measures in ATM to reduce flight emissions	*See Chapter 3*	*See Chapter 3*	I. Incorporate measure of flight emissions (from flight detours) in capacity planning for ATM services II. Demonstrating in simulation study on real flight network that centralized capacity planning and trajectory-independent ATM charges can substantially reduce flight detours and respective emissions

Figure 1.1: Structure of the dissertation

ATM. A motivation for the topic is provided based on the network performance in European air traffic in the past years. A comprehensive literature review on pricing in ATM shows that until today researchers have focused mostly on differentiated, or peak-load, charging to steer demand for ATM services. Therefore, a mathematical framework is developed that allows to evaluate different demand management mechanisms, and forms the basis also of the second and third paper. Furthermore, an efficient heuristic is proposed for solving the $\mathcal{NP}$-hard routing problem, which is used to determine dynamic prices for so-called flexible trajectory products. The performance of the heuristic is compared against that of the exact approach in terms of both solution time and quality. Finally, the developed methodology is applied to test different demand management settings on a small scale case study.

Chapter 3 encompasses the second paper on the topic of capacity management in the strategic phase of ATM. It addresses research question 1 by introducing stochastic capacity decisions to ATM. As a central part of the paper, a concept for a more flexible provision of capacities (i.e., capacity sharing) is presented and analyzed. The introduction includes a brief summary of experts and researchers that recommend 1) a more central role for capacity planning within European ATM, and 2) cross-border capacity sharing to adjust capacities to demand more flexibly. The literature review discusses existing capacity sharing concepts in ATM and other fields. As a problem statement, a two-stage newsvendor problem is

formulated for the strategic capacity planning problem, and adjusted to the setting where capacities can be shared across borders. A stochastic approach for determining strategic capacities is developed which, in contrast to existing models, allows to analyze the benefits of capacity sharing. In the numerical experiments, the stochastic approach is compared against a deterministic benchmark on a large scale case study of 4,000 flights across Western European airspace. Furthermore, the potential of various capacity sharing concepts is analyzed to guide decision-makers for a future ATM architecture in Europe.

Chapter 4 presents the third paper on the environmental benefits of demand and capacity management in ATM. It addresses research question 3 by proposing one demand management and one capacity management measure to reduce flight emissions in Europe. After motivating the focus on reducing aviation emissions, the literature review briefly discusses existing research on sustainable ATM. The proposed demand-capacity balancing measures (network-oriented capacity management and trajectory-independent airspace charging) are then presented and differentiated based on whether the NM or the AUs retain decision making power. Finally, both measures are evaluated on the large scale case study used for the second paper, and implications for capacity and demand management in Europe are discussed.

Chapter 5 concludes the dissertation with a summary of the contributions, critical reflections and recommendations for future research.

1.3 Publications and presentations

The papers developed as part of this dissertation have been published in renowned international journals, and the key concepts have been presented in various occasions to transportation researchers and ATM practitioners. The paper "The value of flexible flight-to-route assignments in pre-tactical air traffic management" was published in *Transportation Research Part B: Methodological* in May 2022 (Künnen and Strauss, 2022). The methodology and findings of the papers were also presented during the 8th INFORMS Transportation Science and Logistics Society Workshop (July 19, 2021), and the World ATM Congress 2021 in Madrid (October 26, 2021). The routing model in particular was also presented to the division of Network Strategy and Development at Eurocontrol (August 11, 2022) to discuss its application in guiding future regulations in the European airspace. A further presentation is planned during the International Conference on Operations Research 2022 at the Karlsruhe Institute of Technology (September 8, 2022).

The paper "Cross-border capacity planning in air traffic management under uncertainty"

is in the second round of review for publication in *Transportation Science*. The paper was co-authored, next to my first supervisor, by Nikola Ivanov, Radosav Jovanovic, Frank Fichert and Stefano Starita. While I developed the methodology, executed the modeling and wrote the original full draft of the paper, the role of the co-authors was to provide guidance from an economic (Frank Fichert), air traffic management (Nikola Ivanov and Radosav Jovanovic) and modeling (Stefano Starita) perspective. The contributions of all authors is mentioned in the Author Statement, which is explicitly included as part of the paper. The work was presented during the World ATM Congress 2021 in Madrid (October 26, 2021), and a meeting of the Eurocontrol Working Group on Environmental Transparency (May 30, 2022).

Finally, the paper "Leveraging demand-capacity balancing to reduce air traffic emissions and improve overall network performance" is currently under review for publication in *Transportation Research Part A: Policy and Practice*. The paper was co-authored by Nikola Ivanov, Radosav Jovanovic and Frank Fichert, next to my first supervisor. Again, their role was to provide guidance from an economic (Frank Fichert) and air traffic management (Nikola Ivanov and Radosav Jovanovic) perspective, while I developed the methodology and original full draft of the paper. The contributions of all authors is mentioned in the paper's Author Statement. The work was presented during the World ATM Congress 2022 in Madrid (June 21, 2022). A further presentation is planned during the ATRS World Conference 2022 in Antwerp (August 24-27, 2022) and the research workshop on Single European Sky and Resilience in ATM in Sofia (September 15-16, 2022). The findings of all three papers will also be presented to, and discussed with, ATM practitioners during a stakeholder workshop at Eurocontrol in Brussels (September 7, 2022).

Chapter 2

Demand management: Flexible flight-to-route assignments

[1]

[1]Published in May 2022 in *Transportation Research Part B: Methodological* under "The value of flexible flight-to-route assignments in pre-tactical air traffic management" with co-author Arne K. Strauss

Abstract

In European air traffic management, there are various discussions regarding the future role of the network manager (NM): in particular, should their role be strengthened to be able to assign flights to specific trajectories, or should airspace users be allowed to freely choose their preferred trajectory, or something in between? If the latter, how would this work? In this paper, we develop a modelling framework that can be adapted to these different settings so as to quantify their effect on key performance indicators.

We focus on the pre-tactical stage of planning air traffic for a future departure day, meaning that airspace capacity budgets are given and incoming flight intentions need to be offered one or several 'trajectories products' for a (possibly dynamically determined) charge. These trajectory products differ in the amount of flexibility that they provide the NM to route the flight shortly before the time of departure. The idea is to reward greater flexibility of the airspace users with lower charges. The airspace user considers the options of trajectory products offered, and chooses one according to a choice model reflecting their preferences in light of various factors including the dynamic trajectory product charges. Shortly before the departure day, the NM decides simultaneously on the routing (within the limits defined by the purchased trajectory products) and on each airspace's sector opening scheme (within the limits of the fixed capacity budgets) so as to minimize the total displacement costs. Charges are set so as to just recover the exogenous cost of capacity budgets and to influence airspace users in their trajectory product choice.

Methodologically, the problem deviates from typical dynamic pricing problems in various major ways (such as featuring a hard boundary condition as well as fairness and revenue neutrality constraints). The problem is cast in the form of a dynamic program with a boundary condition that we show to be NP-hard. We exploit a certain structure in this boundary problem so as to formulate an efficient heuristic. Based on a numerical case study, we find that the use of these trajectory products along with dynamic pricing can be highly beneficial to the extent of achieving a cost performance close to the one obtained if the NM has a mandate to simply assign flights to trajectories shortly before departure. Therefore, this seems an attractive design for the role of the NM, giving airspace users some choice whilst achieving low overall costs.

Keywords: dynamic pricing; demand management; trajectory optimization; route choice

2.1 Introduction

The European ATM environment features significant demand-capacity imbalances leading to costly consequences. According to Eurocontrol (2018), in a typical week in June 2017, demand for ATM services in Europe exceeded capacity 7% of the time (leading to potential delaying of flights), while the sector load was below 60% of capacity half of the time (creating large spare capacity). As a result, 3.6% of flights in the area were affected by ATM-related delays, creating delay costs of EUR 550 million that year. The observed demand-capacity imbalances are mainly due to fragmented and often inflexible capacity planning whilst facing uncertainty in demand (especially non-scheduled flights that account for about 20% of all flights and disruptions in capacity provision, see Eurocontrol (2018)). Once strategic capacity budgets have been determined months in advance, there are only limited options to adjust capacity on short notice. One such option is to change the planned configurations of airspaces subject to the fixed strategic capacity budget. Otherwise, at the pre-tactical planning level, the network manager (NM) can only dynamically implement demand management measures such as assigning delays or to re-route flights so as to minimize the associated costs.

The Wise Persons Group (WPG), which was set up to provide direction for the future of European ATM, states that "if efforts to accommodate demand are not successful and airspace congestion continues, not only would this have a detrimental effect on passengers and other stakeholders, it would also inevitably result in longer flight trajectories, and consequently higher fuel consumption and levels of CO2 emissions", see Wise Persons Group (2019). To counter the inefficiencies, the WPG recommends a stronger role of the NM in routing decisions and "relying on a market-driven approach wherever possible" (Wise Persons Group (2019)).

A natural question arising in this context is how this future role of the NM should look like. In this paper, we focus specifically on developing a modelling framework that can be used to assess the implications of different roles of the NM pre-tactical planning phase. We always assume that the NM is able to decide on which airspace runs in which configuration subject to an exogenously given capacity budget; our focus is on demand-side interactions. From a cost minimization point of view, an extreme setting would see the NM empowered to assign trajectories (including delay or re-routing tasks) shortly prior to departure to all flights. However, it is unlikely to happen in practice since airspace users (AU) would typically also want some influence on their trajectories; in another extreme setting, AUs would have the choice among all feasible trajectories – which can be expected to result in poor cost performance for the system as a whole.

We show that the concept of dynamically priced flexible "trajectory products" (first proposed by Ivanov et al. (2019)) has the potential of offering an excellent trade-off between the two extremes, with some choice being granted to the AUs according to their preferences whilst resulting in an overall cost performance close to that of the empowered NM setting. A flexible trajectory product gives the NM the right to assign the flight at short notice to any trajectory within specified margins of spatial and/or temporal deviation from the Great Circle (shortest path). The larger these margins, the lower the charge for the trajectory product. Charges depend only on the specifics of the trajectory product, not on the actual route taken (origin-destination charging rather than sector-based unit rates), and AUs may choose between multiple trajectory products with different margins and associated charges for a given flight.

The considered problem is the following: the booking horizon of the pre-tactical planning phase starts around 6 months prior to the departure day and ends on departure day. At any time within this period, the AUs can submit their flight plans, in which they request flying from a certain origin to destination at a desired departure time. Once a flight plan gets submitted, the NM needs to decide on the charges to offer to the AU for every trajectory product. The trajectory products differ in the flexibility with which NM is allowed to route a flight. Confronted with these options, the AU then chooses their preferred product and the NM is committed to honouring the trajectory margins set with the purchased product. At the end of the booking horizon, when all flight plans have been submitted and corresponding products have been purchased, the NM needs to decide how to route each flight (in line with the trajectory products) to minimize overall displacement costs. Furthermore, by the end of the booking horizon, the NM needs to have collected (just) sufficient charges on ATM services so as to recover total capacity cost incurred. We call the problem of pricing these trajectory products the "dynamic trajectory pricing problem (DTPP)".

Structurally, the DTPP differs from typical dynamic pricing problems in several fundamental ways: First of all, every flight plan submission must purchase one of the offered options – this leads to different dynamics than in the standard setting where customers may leave without purchasing. In particular, this means that fairness needs to be considered. Secondly, we aim for revenue neutrality meaning that collected charges should closely match the exogenously given capacity cost; overall revenues should neither be substantially larger nor smaller than the fixed capacity cost. Thirdly, around 80% of flights are scheduled and therefore the majority of flight plan arrivals is known to occur at some point; only their precise arrival time and the AUs choice of trajectory products is unknown. Finally, and most

importantly, the DTPP has a hard boundary condition in the form of a routing problem that makes an optimal solution for even moderately-sized instances intractable.

Our main methodological contributions consist of proposing a dynamic programming formulation for the dynamic trajectory pricing problem (DTPP) incorporating fairness and revenue neutrality considerations, and showing that the boundary condition is an NP-hard optimization problem. We provide an efficient heuristic solution approach for the DTPP that can be implemented so as to make dynamic pricing decisions in real-time. Our numerical study provides managerial insights to the debate on the future role of the NM, specifically that dynamically priced flexible trajectory products can achieve an excellent trade-off between AU choice and overall cost minimization.

The paper is organized as follows: In §2.2 we review current literature on pre-tactical ATM and related fields. In §2.3, we provide a formal description of the DTPP, and §2.4 presents an efficient method to solve the problem for realistic instances. We evaluate dynamic pricing policies based on the proposed method in §2.5 and close with recommendations for future research in §2.6.

2.2 Literature Review

From an operations research perspective, the field of air traffic management is generally concerned with the capacity and demand management actions that optimize the flow of air traffic through the network – on a strategic, pre-tactical and operational level. A comprehensive review is given in Barnhart et al. (2012). Within this spectrum of problems, the majority of research until today has focused on the operational demand management actions that optimize routings on the day of operations (known as air traffic flow management, see Mukherjee and Hansen (2009)).

The potential of using differentiated prices to manage demand in pre-tactical ATM has been addressed in a few studies. Castelli et al. (2013) analyse the optimal sector charges the NM should set to maximize their revenues. The authors find in a small real-world test that enroute charges can be an effective instrument to influence the route choice of AU. Steer Davies Gleave (2015) investigate modulation of charges in European airspace and recommend route prices to be set iteratively rather than once at a specific point in time. Finally, Xu et al. (2020) find that a stronger collaboration between AUs and the NM in the pre-tactical phase can significantly reduce delays and detours.

Based on these insights, various researchers have investigated differentiated pricing op-

tions in pre-tactical ATM. A discount/surcharge pricing scheme for managing demand of airspace is investigated by Ranieri and Castelli (2008). The authors compare two options to incentivize AUs to avoid congested sectors: a surcharge to use sectors with high traffic volume or discounts for flights that decide to reroute. In contrast to our approach, they largely use the existing route-charging and only incrementally adjust prices to reflect traffic flow. Jovanović et al. (2014) combine both rerouting and delay incentives and propose discounts/surcharges on sectors that minimize total network cost. The pricing scheme is revenue-neutral, i.e., the revenues only cover capacity provision cost. They use bi-level programming to determine surcharges on congested sectors which in turn subsidise discounts placed on underutilized segments. Their model assumes that demand is known in advance and determine discounts and surcharges accordingly (once per year), while we dynamically adjust charges to demand materializing over time. Furthermore, they only determine binary charges for each sector (peak and off-peak) and therefore offer much less flexibility than the differentiated charges we propose. Bolić et al. (2017) similarly use a bi-level mixed integer programming approach to determine centralised peak-load prices (CPLP). They model the problem as a Stackelberg game: in the first stage the charges are set such that delays and reroutings are minimized, and in the second stage, the AUs seek the cheapest routes with regards to total cost.

Since the IP formulation does not scale to industry-sized problems, they propose in Castelli et al. (2015) two heuristic approaches to solve the CPLP. They find that traffic distribution (in terms of sector load) can significantly be improved through these en-route charges. In contrast to our approach, the CPLP does not explicitly consider customer choice, nor does it anticipate demand over time. The prices set are solely dependent on the capacity usage at the time of booking, while we develop a dynamic model, where charges are adjusted to demand over time. Jovanović et al. (2015) propose a "Reward Predicatbility" model that incentivizes AU to submit flight plans earlier in the process to reduce uncertainty and improve network performance. They effectively adjusts the sector charges whenever a capacity limit is reached, thereby inducing increased charges over time.

The above mentioned models differ from our approach in three important ways: 1) They determine differentiated sector charges rather than dynamic route prices, 2) they do not estimate opportunity cost to set these charges, 3) they assume demand is deterministic and known in advance. Thus, to the best of our knowledge, no research has explored the demand management of ATM services in the pre-tactical phase via dynamic trajectory prices, incorporating opportunity cost and stochastic demand. The idea to strengthen the role of

the NM in Europe by introducing flexible products is first discussed in Ivanov et al. (2019).

Our work is related to papers investigating optimal capacities in the strategic phase in as far as the strategic capacity budget is an exogeneous input to our model. Starita et al. (2020) provide such a strategic capacity planning model; we use a similar formulation for the cost estimation as theirs, but solve it heuristically in a different (faster) way since we need to make much faster-paced dynamic decisions. Outside the context of ATM, the problem of estimating opportunity cost in the presence of a hard routing problem is investigated in dynamic tolling for traffic networks (see Rambha and Boyles (2016)), as well as dynamic pricing in attended home delivery (AHD, see Yang and Strauss (2017) Yang et al. (2016)). The latter work has inspired our approach in that we also design a "foresight policy" that attempts to anticipate future demand when estimating opportunity cost via insertion costs.

2.3 Problem Statement

In this section, we present the mathematical modeling framework of the dynamic trajectory pricing problem (DTPP) in pre-tactical ATM. The formulation is kept sufficiently general to accommodate the three settings that we seek to investigate, namely either full flexibility to assign flights by the NM, full choice of trajectories by the AUs, or a mix in the form of AU choice between flexible trajectory products.

2.3.1 Problem definition and notation

In all settings that we investigate, the AU requests the ATM service for a certain flight and the NM sets the service charge; however, the options for the AU vary in each setting. We always plan for a single day of departure. The pre-tactical ATM process starts at a fixed number of days prior to the departure day and ends on departure day. In particular, we consider a discrete booking horizon from $t = 1, \dots, T$, where T is the cut-off time after which no bookings are permitted. Within the booking horizon, any AU can submit a flight plan for a flight $f \in \mathcal{F}$, in which they request flying from some origin to some destination at a certain departure time. To incorporate non-scheduled flights, for which no such information exists in advance, set $\mathcal{F}$ includes all potential combinations of origins, destinations and departure time. This can be achieved by pooling historic origin-destination pairs and by discretizing departure time into small time intervals (e.g., 5 minutes). At any time t, we denote the set of (scheduled) flights for which flight plans have already been submitted by F_t (F_t^G), and their complements by $\bar{F}_t$ ($\bar{F}_t^G$). Note that $F_t^G \subseteq F_t$ and $F_t \cup \bar{F}_t = \mathcal{F}$. Booking requests

for non-scheduled flights arrive according to a Poisson process with λ^N, while arrivals for scheduled flights are modeled by a pure "death process" with λ^G (see §2.3.1). Time t is chosen sufficiently small so that at most one request arrives per time period.

We introduce new trajectory product types that the AU can book, which determine how flexibly the NM can decide to route the flight (see §2.3.1). Once the flight plan is submitted for a certain flight f and day, the NM needs to decide on the price vector p^f, with prices for each product type $z \in Z$, to charge to the AU. Confronted with prices $p^f = (p_z^f)_{z \in Z}$, the AU chooses product type z with probability $P_z(p^f)$. Based on the purchased product types, the NM then needs to decide through which sectors to route each flight to keep network cost low. Section 2.3.1 describes in detail the modeling of the DTPP as a Markov decision process.

Definition of products

As mentioned above, the product type determines the flexibility conditions under which a flight can be routed on the day of operation. These are relevant in the setting in which the AU chooses between different trajectory products. We define a route as the spatio-temporal trajectory of a flight. Let the day of operations be divided uniformly into discrete operating periods u in $U = \{1, \dots, U^{max}\}$. Any flight f can then be routed through routes $r \in R^f$, where each r is a sequence of elementary sector- and time-combinations. (Bolić et al. (2017) estimate that a typical flight in Europe only chooses among 4 clearly distinct routes, so that we can assume a finite route set.) Any route $r \in R^f$ comes with flight displacement cost d_r^f which reflect the additional fuel and delay costs generated by routing a flight through r, relative to the shortest distance and no delay.

With this notation, we define n product types in set $Z = \{1, \dots, n\}$. If an AU buys product type $z \in Z$ for flight f, the NM commits to routing the flight through $R_z^f \subseteq R^f$. The flexibility with which flights can be routed increases with z so that we have $R_1^f \subset \cdots \subset R_n^f = R^f$ for each flight f, where R_1^f are those routes that are operationally close to the flight's Great Circle Distance (GCD). The GCD describes the shortest distance between any two points on the surface of a sphere (i.e., the earth), and in this case represents the shortest possible route between any two city pairs. Since the price of a product depends on product type z and flight f for which it is purchased, we define a product j as any combination (f, z).

State space, action space and transition function

In our decision model, the NM uses the latest booking information (state space) to price the trajectory products (action space), and then observes the response by the AU (transition

function) to conclude the booking request.

State space The state space needs to contain all information that we require to take an action at booking time t. The state of the DTPP is fully described by $X_t \in \chi \subset \mathbb{N}^{|\mathcal{F}|+1}$, which contains in positions $1, \ldots, |\mathcal{F}|$ the product type that has been purchased until t (if any) and the prices p^f that were offered, for each flight $f \in \mathcal{F}$. In particular, the k-th element $X_{t,k}$ for $k = 1, \ldots, |\mathcal{F}|$ is defined as:

$$X_{t,k} = \begin{cases} 0 & \text{if flight plan has not been submitted} \\ (z, p^f) & \text{if product type } z \text{ was purchased by flight } k \text{ at price offer } p^f. \end{cases}$$

An exemplary state vector ith two product types is illustrated below, where n_G and n_N is the number of scheduled and non-scheduled flights, respectively:

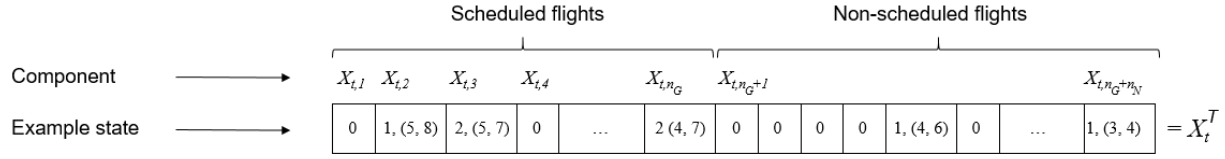

Figure 2.1: Example of a DTPP state space vector

Action space If we are in state X_t and a flight plan is submitted for flight f, we need to decide on prices to offer to the AU for every product type $z \in Z$. According to common business practice, a limited number of discrete price points is suitable. Therefore, we develop a vector with discrete, relative price points that can be applied to all flights. For that purpose, we define a benchmark price $r\bar{e}v^f$ for every origin-destination pair reflecting the revenues needed to cover capacity provision cost. To compute benchmark prices for each flight, we proceed in three steps: First, we use historic flight patterns to determine the average share of flights for each combination of origin-destination pair and aircraft type (flying on these pairs). Second, we define a relative cost index for each combination of origin-destination pair and aircraft type. This cost index shows the relative cost generated by one such combination over another, and reflects that longer flights (and larger aircraft) cause higher cost than shorter flights (and smaller aircraft). Lastly, we split the total capacity provision cost among all combinations of origin-destination pair and aircraft type according to their relative share and cost index. Given parameters $r\bar{e}v^f$, each price point rev_z^f can then be represented as the

percentage p_z^f of the flight's benchmark price that is charged to the AU, i.e.,

$$rev_z^f = p_z^f r\bar{e}v^f, \qquad f \in \mathcal{F}, z \in Z, t = 1, \ldots, T. \tag{2.1}$$

With this notation, we can replace the action space of pricing vector $(rev_z^f)_{z \in Z}$ with the action space of $p^f = (p_z^f)_{z \in Z}$ for each flight, where p_z^f is chosen from a finite set of scaling factors $Pr = \{p_i : i = 1, \ldots, I\}$. For instance, a price p_z^f of 1.1 means that we are charging 10% more than the benchmark price $r\bar{e}v^f$. Since we are pricing up to n product types for each flight, our action space has complexity I^n at each time. We always need to offer at least one product type because we cannot deny the service offering. In fact, we always offer all product types in order to maximize the choice for AUs.

Transition function Having decided on pricing offer p^f at time t, the transition from state X_t to X_{t+1} depends on the customer choice outcome as well as arrival rates of upcoming flights. The product choice by the AU is governed by a choice model; it defines the probability $P_z(p^f)$ that an AU purchases product type z if confronted with pricing offer p^f. Note that, in contrast to traditional revenue management problems, the AU *has* to choose one of the products offered in the booking process (*booking obligation*), i.e.,

$$\sum_{z \in Z} P_z(p^f) = 1, \quad \forall f \in \mathcal{F}.$$

This condition is particular to the DTPP and requires us to impose further constraints: To ensure that the pricing mechanism does not abuse the booking obligation by always setting maximum prices, we implement a revenue neutrality and fairness requirement, which is discussed in §2.3.2.

Arrival process

In contrast to traditional revenue management problems, we know almost certainly that a large share of customers (i.e., scheduled flights) will eventually "arrive" to the booking process; we just do not know when. To model the arrival process, we first require that once a flight (scheduled or non-scheduled) has entered the booking process, it does not arrive again. Let $\lambda_t^f(F_t)$ be the arrival probability of a particular flight f at t, given that flights F_t have already arrived so far. We have:

$$\lambda_{t'}^f(F_t) = 0, \quad \forall f \in F_t, t' = t, \ldots, T.$$

In the following, we therefore focus on arrival probabilities of remaining flights $\bar{F}_t$, where we differentiate between scheduled and non-scheduled flights. For non-scheduled flights we can assume that AUs arrive according to a Poisson process with arrival rate λ^N. This arrival rate governs the flight arrival event itself; the flight specifics (origin, destination and departure time) are uniformly sampled from a large finite set. For scheduled flights, the arrival process is more complex because the arrival rate at t depends on the remaining population $\bar{F}_t^G$. Since we know that most (if not all) scheduled flights will enter the booking process at some time until T, we expect a higher arrival rate from t to T if few scheduled flights have arrived until t. In particular, we require:

$$\sum_{t'=t}^{T} \lambda_{t'}^f(F_t) \approx 1, \quad \forall f \in \bar{F}_t^G.$$

In Parlar et al. (2018), the authors discuss a similar setting when modeling the arrival of customers to exclusive-use airline check-in counters, where customers can only use certain counters to check in for their flight. As in our setting, the authors expect most (if not all) passengers of a flight to arrive to check-in before the counter closes. They model the arrival as a pure "death process", where the time until arrival of customers is exponentially distributed with parameter λ_t^G. We assume that the arrival time distribution in the pre-tactical ATM process can also be reasonably approximated as exponential since the AUs are incentivized to submit their flight plans early in the process to secure attractive trajectory options. Parameter λ_t^G can be interpreted as the probability that a certain scheduled flight arrives within the next time period after t.

In Parlar et al. (2018), the authors estimate λ_t^G based on historic arrival patterns. Let $\tau_0 < \tau_1 < \tau_2 < \cdots < \tau_m$ be the points in time (i.e., epochs) during which individual arrivals of scheduled flights occured in one such historic arrival pattern, and let $q(t)$ specify the last epoch before time t. Also, let $x_0, x_1, \ldots, x_m$ be the observed number of scheduled flights that have not yet arrived to the booking process at the start of each epoch, where $x_m > 0$ in case of cancellations. Then the cumulative time until arrival can be expressed as $D_t = \sum_{i=q(t)}^{m-1} x_i(\tau_{i+1} - \tau_i) + x_m(T - \tau_m)$. Specifically, D_t is the expected number of time intervals that we need to wait until an arrival of any flight $f \in \bar{F}_t^G$ occurs; it will be estimated as an average over multiple historic arrival patterns. The probability of arrival at any time t can then be estimated by $\lambda_t^G = 1/D_t$, which will be updated dynamically after every arrival. In summary, we can estimate the arrival probability for any scheduled or non-scheduled flight

as:

$$\lambda_t^f(\bar{F}_t) = \begin{cases} \lambda_t^G = \frac{1}{D_t} & \text{for all } f \in \bar{F}_t^G \\ \lambda^N & \text{for all } f \in \bar{F}_t \backslash \bar{F}_t^G \\ 0 & \text{for all } f \in F_t. \end{cases}$$

2.3.2 Dynamic programming formulation

Value function

Our goal is to determine a policy that determines price parameters $p^f(X_t)$ to offer for any flight f, given existing bookings X_t at time t, such that the AU is steered towards a product type with lowest expected displacement cost (including delay and rerouting cost). Henceforth we will omit subscript t in X_t since the time will follow immediately from the dynamic program recursion. The policy needs to consider at any time t the set $\bar{F}_t$ of potential flights for which no trajectory product has been purchased yet. Let $V_t(X)$ be the value of being at state X at time t, meaning the minimum expected cost we need to bear from t until cut-off time $T+1$, given product purchases in X.

The value function is then defined by:

$$\begin{aligned} V_t(X) &= \sum_{f \in \bar{F}_t} \lambda_t^f(\bar{F}_t) \min_{p^f} \left\{ \sum_{z \in Z} P_z(p^f) V_{t+1}(X \cup (f,z,p^f)) \right\} + \left[1 - \sum_{f \in \bar{F}_t} \lambda_t^f(\bar{F}_t) \right] V_{t+1}(X) \\ &= \sum_{f \in \bar{F}_t} \lambda_t^f(\bar{F}_t) \min_{p^f} \left\{ \sum_{z \in Z} P_z(p^f) \left[V_{t+1}(X \cup (f,z,p^f)) - V_{t+1}(X) \right] \right\} + V_{t+1}(X), \quad X \in \chi. \end{aligned} \tag{2.2}$$

where $V_{t+1}(X \cup (f,z,p^f) - V_{t+1}(X) =: \Delta_{(j,p)} V_{t+1}(X)$ is the opportunity cost of selling product $j = (f,z)$ with price set p^f at state X_t. At any time t, a request for flight f arrives with probability $\lambda_t^f(F_t)$ and the AU decides to purchase product type z given pricing offer p^f with probability $P_z(p^f)$. In that case we incur the expected cost of state $X \cup (f,z,p^f)$ we are moving to. If no request arrives at time t, we remain in state X for $t+1$.

Boundary condition At the end of the booking horizon we need to determine the minimal displacement cost $V_{T+1}(X)$. For this, we denote by $D(X,W,H)$ an oracle that provides the minimum displacement cost for routing flights via trajectory options defined by state X, given capacity budget vector H and sector capacity uncertainty W. Capacity budget $H = (H_a)_{a \in A}$

specifies the total sector hours that an airspace a can use, and is set at the strategic phase. Parameter W governs the uncertainty around actual sector capacities at $T+1$, which are subject to operational fluctuations, e.g. employee absence or weather conditions.

As mentioned before, to control the pricing behavior, there are two further soft considerations that we need to include in the boundary condition: Firstly, the total revenues generated by the pricing policy need to be in line with total capacity provision cost; that is, revenues need to be large enough to recover these cost, and can only exceed them by a certain margin (*revenue neutrality requirement*). Secondly, the range of prices offered to an AU for different product types should not be excessively large; in particular, the range should reflect the difference in opportunity cost between the product types (*fairness requirement*). While revenue neutrality regulates the total sum of charges, the fairness requirement regulates the range of charges among product types. Let θ^{RN} and θ^{FR} be the penalties for violating the revenue neutrality and fairness requirement, respectively. Furthermore, we denote by $\hat{p}^f = \sum_z P_z(p^f) p_z^f$ the expected price paid for flight f and by $Var(\cdot)$ the variance of any set.

Then the boundary condition including soft considerations is given by:

$$V_{T+1}(X) = \mathbb{E}_{X,W}[D(X,W,H)] + \theta^{RN}\epsilon^{RN}(X) + \theta^{FR}\epsilon^{FR}(X), \quad X \in \chi, \tag{2.3}$$

where

$$\epsilon^{RN}(X) := \sum_{f \in F_T} |1 - \hat{p}^f|, \tag{2.4a}$$

$$\epsilon^{FR}(X) := \sum_{f \in F_T} Var(p^f). \tag{2.4b}$$

Definition (2.4a) sets ϵ^{RN} as the absolute deviation between the average collected price and cost-neutral price of 1, respectively. Multiplying both parts of the subtraction with benchmark prices rev^f for each flight would result in the deviation between capacity budget cost and collected revenues, which is what we require to determine revenue neutrality. Since we do not know in advance how many flights will arrive, it is possible that total revenues exceed capacity cost if unexpectedly many flights arrive, or fall short of capacity cost if unexpectedly few flights arrive. However, across multiple operating days, these deviations will cancel out so that we ensure revenue neutrality over time. Definition (2.4b) defines ϵ^{FR} as the variance among prices offered for flight f. This way we prevent the optimization from setting excessively large price differences between product types and thereby implicitly imposing a

product on the AU. We will only charge largely different prices if the difference in opportunity cost between products outweighs the penalty associated with the fairness condition. Both components ϵ^{RN} and ϵ^{FR} are then used with penalties θ in boundary condition (2.3) to model these soft considerations. The penalty terms should be set according to the application at hand, and reflect the choice model $P_z(p^f)$. Even if definitions (2.4a) and (2.4b) are constructively similar, they cannot be reduced to one condition because (2.4a) controls the average price while (2.4b) controls its variation.

Joint sector-opening and routing optimisation

Computing $D(X, W, H)$ exactly requires solving a problem that determines a) the optimal trajectories for each flight and b) the sectors to open at each operating time u (the so-called sector-opening scheme), given bookings in X, uncertainty W and capacity budget H, such that total displacement cost are minimized. In evaluating $D(X, W, H)$, we assume the expected displacement cost d_r^f for routing f through r to be given.

The following model is based on Starita et al. (2020). Let x_r^f be the decision whether f is routed through r, and let y_{acu} be the decision whether airspace a operates in configuration $c \in C^a$ at operating time u. A configuration $c \in C^a$ specifies a partitioning of airspace a into (elementary or collapsed) sectors l, which in turn are stored in L^c. The set of elementary sectors e that form any sector l are denoted by E_l. Operating time intervals are chosen sufficiently large (e.g, 1 hour) so that configurations can be changed between intervals. Parameter $\bar{h}_{ac}$ represents the sector-time units consumed if airspace a operates in c. The indicator b_{freu} equals 1 if flight f on route r uses sector e at time u, and 0 otherwise. For every $f \in F_t$ we define z as the product type that was purchased by the AU, so that R_z^f are the routing options for flight f. Finally, sector capacity $\mathcal{K}_l$ specifies how many flights can enter sector l within one time interval, based on possibly reduced capacities W. The integrated routing

and sector opening problem (IRSOP) is given as:

$$\textbf{IRSOP: } D(X_t, W, H) = \min_{x,y} \sum_{f \in F_t} \sum_{r \in R_z^f} d_r^f x_r^f \tag{2.5a}$$

$$\text{s.t.} \sum_{u} \sum_{c \in C^a} \bar{h}_{ac} y_{acu} \leq H_a \qquad a \in A \tag{2.5b}$$

$$\sum_{f \in F_t} \sum_{r \in R_z^f} \sum_{e \in E^p} b_{freu} x_r^f y_{acu} \leq \mathcal{K}_l \quad a \in A, c \in C^a, l \in L^c, u \in U \tag{2.5c}$$

$$\sum_{c \in C^a} y_{acu} = 1 \qquad a \in A, u \in U \tag{2.5d}$$

$$\sum_{r \in R_z^f} x_r^f = 1 \qquad f \in F \tag{2.5e}$$

$$x_r^f \in \{0,1\} \qquad f \in F, r \in R_z^f \tag{2.5f}$$

$$y_{acu} \in \{0,1\} \qquad a \in A, c \in C^a, u \in U. \tag{2.5g}$$

The objective function minimizes total flight displacement cost over all flights and routes. Constraint (2.5b) defines the feasible configurations based on the capacity budget H_a for each airspace. Constraint (2.5c) ensures that the sector capacity is not exceeded: If we operate under configuration c at u, we restrict the capacity of any sector in L^c to $\mathcal{K}_l$; else the left-hand side reduces to 0 and leviates the restriction. Constraint (2.5d) ensures that each airspace operates under one configuration at any time, (2.5e) ensures that exactly one route is assigned to each flight, and (2.5f) and (2.5g) model the binary condition for our decision variables.

Real-time control policy

If the dynamic program for the value function is solved, it can be used as input for the online decision policy. As soon as an AU submits a request for a flight f, we need to decide which prices $p^f(X)$ to offer. We see from the dynamic program in (2.2) that this online decision can be made as follows:

$$p^f(X) = \arg\min_{p^f} \sum_{z \in Z} P_z(p^f) \Delta_{(j,p)} V_{t+1}(X). \tag{2.6}$$

We set prices p_t^f for flight f such that expected opportunity cost $\Delta_{(j,p)}V_{t+1}(X)$ is minimized. The pricing policy is dynamic in that we adjust the price vector p^f based on incoming bookings over time. Determining optimal prices in (2.6), given known value function $V_t(X)$, represents an assortment optimization problem whose solution technique depends on the customer choice model describing $P_z(p^f)$. Note that due to the assumption of finite price points, we need to choose one price vector from a finite set of possible vectors. Since our action space has complexity I^n, we can in fact fully enumerate the price vectors and their respective objective value in (2.6) if the number of price points I is small. For large I, one of the assortment optimization techniques in Strauss et al. (2018) can be applied.

In summary, the key to making optimal decisions is to quantify the opportunity cost, which in turn depends on the value function. Since the latter is intractable for realistic problem sizes, we need to find high quality approximations of the opportunity cost. This is the subject of the following section, where we discuss approaches for approximating opportunity cost offline so that they can be used in the real-time decision policy.

2.4 Approximation of opportunity cost

As explained above, we require an approximation of the opportunity cost as an input to the pricing policy. The opportunity cost associated with selling product $j = (f, z)$ under price offer set p^f represents all future cost implications from this transaction. In particular, it includes displacement and penalty cost implications. Let X_{T+1}^S and W^S represent realizations of the respective uncertainties for a scenario $S \in \mathcal{S}(X)$. In particular, a scenario describes a forecast of flight arrivals and actual capacity for departure day. Here, $\mathcal{S}(X_t)$ denotes the population of scenarios for $X_t \in \chi$; it depends on X_t because bookings that have already been made will always form part of X_{T+1}^S. Note that X_{T+1} represents the realized state at time $T+1$ with known bookings, while X_{T+1}^S represents a forecast of X_{T+1} at any time $t \leq T$ under scenario $S \in \mathcal{S}(X_t)$. Furthermore, let π^S denote the optimal pricing policy determined in hindsight given scenario S, and $X_{T+1}^{\pi^S,S}$ the state under scenario S given that policy π^S is executed. To compute opportunity cost, we use the following value function approximation:

$$V_t(X_t) \approx \hat{V}_t(X_t) := \sum_{S \in \mathcal{S}(X_t)} \min_{\pi^S} \frac{V_{T+1}(X_{T+1}^{\pi^S,S})}{|\mathcal{S}(X_t)|}. \tag{2.7}$$

That means, we approximate the value function at t with the expected value at $T+1$, evaluated over $S \in \mathcal{S}(X_t)$, given that policy π^S is executed. This approximation represents

an asymptotically lower bound on $V_t(X_t)$ for a sufficiently large set of sample scenarios.

Proposition 1. *Let $\hat{\pi}$ denote an arbitrary pricing policy. Then the following holds:*

$$\lim_{|\mathcal{S}|\to\infty} \sum_{S\in\mathcal{S}(X_t)} \min_{\pi^S} \frac{V_{T+1}(X_{T+1}^{\pi^S,S})}{|\mathcal{S}(X_t)|} \leq V_t(X_t) \leq \lim_{|\mathcal{S}|\to\infty} \sum_{S\in\mathcal{S}(X_t)} \frac{V_{T+1}(X_{T+1}^{\hat{\pi},S})}{|\mathcal{S}(X_t)|}.$$

Proof. In general, value $V_t(X_t)$ can be represented as the expected value V_{T+1} at cut-off time resulting from optimal pricing policy π, given the uncertainties X_{T+1} and W (reflected by scenarios S). The lower bound is then deducted as below:

$$V_t(X) = \min_{\pi} \mathbb{E}_S[V_{T+1}(X_{T+1}^{\pi,S})] \underset{SAA}{=} \min_{\pi} \lim_{|\mathcal{S}|\to\infty} \sum_{S\in\mathcal{S}(X_t)} \frac{V_{T+1}(X_{T+1}^{\pi,S})}{|\mathcal{S}(X_t)|} \geq \lim_{|\mathcal{S}|\to\infty} \sum_{S\in\mathcal{S}(X_t)} \min_{\pi^S} \frac{V_{T+1}(X_{T+1}^{\pi^S,S})}{|\mathcal{S}(X_t)|}.$$

We use sample average approximation (SAA) to approximate the expectation over S with the average across these scenarios. The remaining inequality is given because any policy π^S determined in hindsight based on S will deliver a value $V_{T+1}(X_{T+1}^{\pi^S,S})$ at least as low as under policy π. In contrast, to prove the upper bound we use the fact that $V_{T+1}(X_{T+1}^{\pi,S})$ is at least as low as the value under a fixed policy $\hat{\pi}$. Using SAA again shows the proposed relation:

$$V_t(X_t) = \min_{\pi} \mathbb{E}_S[V_{T+1}(X_{T+1}^{\pi,S})] \leq \mathbb{E}_S[V_{T+1}(X_{T+1}^{\hat{\pi},S})] \underset{SAA}{=} \lim_{|\mathcal{S}|\to\infty} \sum_{S\in\mathcal{S}(X_t)} \frac{V_{T+1}(X_{T+1}^{\hat{\pi},S})}{|\mathcal{S}(X_t)|}.$$

□

Using the approximation in (2.7), we obtain for the opportunity cost:

$$\Delta_{(j,p)} V_{t+1}(X_t) \approx \Delta_{(j,p)} \hat{V}_{t+1}(X_t) = \sum_{S\in\mathcal{S}(X_t)} \frac{\Delta_j D(X_{T+1}^S, W^S, H)}{|\mathcal{S}(X_t)|} + \Delta_p \epsilon^{RN} + \Delta_p \epsilon^{FR}. \tag{2.8}$$

In other words, we assume that the opportunity cost of selling product j at price set p^f can be reasonably approximated with the expected insertion cost $\Delta_j D(X_{T+1}^S, W^S, H)$, evaluated over $S \in \mathcal{S}(X)$, alongside the penalties from revenue neutrality and fairness requirements. The quality of the approximation in (2.8) depends on our ability to adequately model the uncertainty involved in our problem setting between $t+1$ and $T+1$. Therefore, we determine $\Delta_{(j,p)} \hat{V}_{t+1}(X)$ in three steps, which are detailed in the following subsections:

1. Use latest information at time t to generate samples $S = (X_{T+1}^S, W^S)$, see §2.4.1;

2. Determine routing of flights with lowest cost $D(X^S_{T+1}, W^S, H)$ for each sample, see §2.4.2;

3. Determine cost $\Delta_j D(X^S_{T+1}, W^S, H)$ as well as penalty parameters $\Delta_p \epsilon^{RN}$ and $\Delta_p \epsilon^{FR}$ for adding product j at price set p^f to the routing of any sample (X^S_{T+1}, W^S), see §2.4.3.

2.4.1 Sampling strategy to model uncertainty

The approach is somewhat similar to the foresight heuristic proposed by Yang et al. (2016) for solving a routing problem in attended home delivery. They used a set of final historic routes to estimate insertion costs. In contrast, in our application we already know more about future demand, specifically scheduled traffic. Hence, we always use the latest information of the current booking process to evaluate the expectation. The information we have on hand at booking time t is two-fold: The flights that have already terminated the booking process along with the products that they have purchased (X_t), and the set of scheduled flights that will still enter the process until cut-off time $(\bar{F}^G_t)$. Therefore, in order to compute expected displacement cost $D(X_{T+1}, W, H)$, we need to address two remaining uncertainties:

- Non-scheduled flights that arrive until cut-off, as well as products that upcoming scheduled and non-scheduled flights purchase (X_{T+1});
- Actual capacity per sector-time unit on departure day (W).

As discussed, we create sample scenarios $S = (X^S_{T+1}, W^S)$ that differ with regard to the two inputs mentioned above. For X^S_{T+1}, we randomly sample non-scheduled flights from a finite set of potential flights, whose origins, destinations and departure times reflect the real population of flights. The number of flights we sample is drawn from a normal distribution (given the large number of flights, we can use a continuous domain) with mean equal to $\mu_F - |F_t| - |\bar{F}^G_t|$, where μ_F is the average expected number of flights on operating day. Furthermore, we assume as part of our policy that all upcoming flights purchase the most flexible product type. This way, we can choose among all $r \in R^f$, letting our model determine the optimal route. The assumption is reasonable because we expect the pricing policy to let AUs choose less flexible products only if this does not impact network performance. Initial experiments confirm that more selective forecasts of product types (e.g., stratified sampling or flight-based inferences) do not improve results. Finally, the realized capacity W^S is modeled by choosing at random one elementary sector in each airspace, and reducing capacity of this sector (and all collapsed sectors containing it) in line with observed historic rates.

2.4.2 Metaheuristic approach for routing problem

In a first step, we need to determine the routing with lowest displacement cost $D(X^S_{T+1}, W^S, H)$ in (2.5a) for each sample scenario (X^S_{T+1}, W^S), which requires solving the IRSOP. However, this problem is $\mathcal{NP}$-hard, as shown below.

Theorem 1. *The IRSOP described by* (2.5a)–(2.5f) *is $\mathcal{NP}$-hard.*

Proof. We show that by fixing certain variables in the integrated routing and sector opening problem, the problem represents an instance of the *multi-choice multidimensional knapsack problem* (MMKP), which is known to be $\mathcal{NP}$-hard Martello (1990). We choose at random a configuration set $C' = \{c'_{au} \in C^a : a \in A, u \in U\}$ with configurations for each airspace and operating time such that condition (2.5b) holds. (If no such configuration set exists, the problem is infeasible.) We set $y_{ac'u} = 1$ for all configurations $c' \in C'$, and $y_{acu} = 0$ otherwise. With this variable fixing, constraints (2.5b),(2.5d)–(2.5g) become redundant, and the remaining problem changes to:

$$
\begin{aligned}
D(X_t, W, H) = \min_x \; & \sum_{f \in F_t} \sum_{r \in R^f_z} d^f_r x^f_r \\
\text{s.t.} \; & \sum_{f \in F_t} \sum_{r \in R^f_z} \sum_{e \in E^l} b_{freu} x^f_r \leq \mathcal{K}_l && c' \in C', l \in L^{c'}, u \in U \\
& \sum_{r \in R^f_z} x^f_r = 1 && f \in F \\
& x^f_r \in \{0,1\} && f \in F, r \in R^f_z.
\end{aligned}
$$

The problem above represents an instance of the MMKP. Within the knapsack-analogy, we need to choose one item (route) from each group (flight) such that total payoff is maximized (displacement cost is minimized), while the capacity limit of the knapsack (airspace) is not exceeded on any dimension (sector-time resource). □

We can also interpret the proof as follows: to determine an exact solution to the integrated routing and sector opening problem we would need to solve an MMKP for every possible combination of airspace configurations, of which there are $\prod_a |C^a|^{U^{max}}$. In order to find an approximate solution in polynomial time, we decouple the sector-opening from the routing procedure: We first determine the best candidate configuration set C' and then solve the remaining routing problem as *one* instance of the MMKP.

Determine best candidate configuration

To find configuration set C', we first assign each flight to the route with lowest displacement cost (giving allocation x), and then determine the feasible configuration that creates the lowest total capacity shortage for routing x. Let parameters k_{acu} represent the capacity shortage (i.e., the number of flights that exceed sector capacity limits) in airspace a, configuration $c \in C^a$ and time unit u. We have $k_{acu} := \sum_{l \in L^c} \left(\sum_{e \in E^l} \sum_{f \in F_t} \sum_{r \in R^f_{z'}} b_{freu} x^f_r - \mathcal{K}_l \right)^+$, where $x^+ := \max\{x, 0\}$. Configuration set C' can then be determined with the following configuration ILP:

$$\textbf{CILP:} \quad \min_y \sum_{a,c,u} k_{acu} y_{acu} \tag{2.9a}$$

$$\text{s.t.} \sum_u \sum_{c \in C^a} \bar{h}_{ac} y_{acu} \leq H_a \qquad a \in A \tag{2.9b}$$

$$\sum_{c \in C^a} y_{acu} = 1 \qquad a \in A, u \in U \tag{2.9c}$$

$$y_{acu} \in \{0, 1\} \qquad a \in A, c \in C^a, u \in U. \tag{2.9d}$$

Configuration set $C' = \{c'_{au} : a \in A, u \in U\}$ consists of individual configurations c'_{au} for which $y_{ac'u} = 1$ for each airspace and operating time. The described program decomposes by airspace. For each airspace, the resulting problem represents a multiple choice knapsack problem (MCKP), which again is $\mathcal{NP}$-hard. For most airspaces, the number of configuration options is sufficiently low so that the problem can still be solved exactly in reasonable time; in any other case, we revert to heuristic approaches for the MCKP (such as Pisinger (1995)). If after solving the CILP there is still spare capacity in an airspace (i.e., constraint (2.9b) is not binding for $a \in A$), we assign the next-higher configuration for a at the operating time u at which shortage $k_{ac'u}$ is currently the highest. We iterate until there is no spare capacity.

Determine best routing within configuration

In a second step, we apply an MMKP heuristic based on the approach by Moser et al. (1997) to determine the optimal routing of flights given airspace configuration set C'. The approach is summarized in Algorithm 1. We establish an initial infeasible solution by assigning each flight to the route with lowest displacement cost. In each iteration we then reassign flights on the most-violated sector until all sectors are within capacity limits $\mathcal{K}_l$. Let $L' = \{l \in L^{c'} : c' \in C'\}$ represent the sectors defined by configuration C'. Based on the current infeasible routing, we determine the most-violated sector $l^* \in L'$ as the one with the largest relative capacity

shortage $\bar{k}_l$. Let $w_{frl} = \sum_{e \in E^l} b_{freu}/\mathcal{K}_l$ be the relative "weight" of flight f and route r on sector $l \in L'$, and let r' be the currently selected route for f. Then $\bar{k}_l = \sum_{f \in F_t} w_{fr'l}$ for each sector $l \in L'$.

To decide which flight to reassign to another route, we first determine flights and routes with a change of capacity usage on l^*, i.e. with positive value $w_{fr'l^*} - w_{frl^*}$. For these flights and routes, we then compute decision parameter γ_r^f that weighs the change in displacement cost against the change in capacity, as follows,

$$\gamma_r^f = \frac{d_r^f - d_{r'}^f - \sum_{l \in L'} \mu_l (w_{fr'l} - w_{frl})}{w_{fr'l^*} - w_{frl^*}},$$

where μ_l is the Lagrange multiplier for sector $l \in L'$. The general idea is to iteratively approximate with μ_l the dual value of the constraint on sector l; for a detailed background of the heuristic see Moser et al. (1997). We then reassign the flight with the lowest ratio γ_r^f from r' to r, update the Lagrange multiplier μ_{l^*} (as shown in Algorithm 1), and reiterate the process until a feasible solution is found.

Algorithm 1 MMKP-heuristic for ATM routing model

Input: Set of flights F and product types Z, candidate best configuration C'
1: **Initialize:** $r'_f := \arg\min_{r \in R^f} d_r^f$ for all $f \in F_t$, $\mu_l := 0$ for $l \in L'$
2: **Establish feasible solution:** Iterate until $\bar{k}_l \leq 1 \forall l \in L'$
3: Compute $\bar{k}_l$ and set sector $l^* := \arg\max_l \bar{k}_l$
4: For flights with positive $w_{fr'l^*} - w_{frl^*}$ on l^*, store route with lowest γ_r^f.
5: Determine flight f and route r with lowest γ_r^f, reassign flight and update Lagrange Multiplier: $r' = r, \mu_{alu^*} = \mu_{alu^*} + \gamma_r^f$
6: **Improve feasible solution:** Iterate until no further improvement found
7: Compute δ_r^f for all $f \in F_t$, $r \in R^f$.
8: Find flight and route with largest δ_r^f and set $r'_f := r$.
Output: Routing $R^* = \{r'_f : f \in F\}$

Finally, we improve the feasible solution further by using potential spare capacity on sectors with non-binding constraints. To do that, we compute an improvement factor δ_r^f for all routes that are currently not selected (i.e., $r \neq r'$) on each flight:

$$\delta_r^f = d_{r'}^f - d_r^f \quad \text{if } d_{r'}^f > d_r^f \text{ and } \bar{k}_l - w_{fr'l} + w_{frl} \leq 1, l \in L'.$$

The flight f and route r with the largest value δ_r^f is then reassigned from route r' to r, and the procedure terminates with final routing $R^* = \{r'_f : f \in F\}$ when no further improvement

is found (i.e., $\delta = \emptyset$). As shown in Moser et al. (1997), Algorithm 1 has complexity $\mathcal{O}(m(n-g)^2 + mn)$, where $m = |L'|$ is total number of sectors given configuration C', $n = \sum_{f \in F_t} R^f$ is total number of flight-route combinations, and $g = |F_t|$ is total number of flights.

The heuristic by Moser et al. (1997) was chosen because it is one of few MMKP heuristics that start from a very good (i.e., low-cost) but infeasible solution and iteratively establish feasibility, rather than vice versa. In European ATM, many flights may be assigned to their shortest route (having least displacement cost) without violating capacity constraints, so that we expect to require fewer iterations to establish feasibility from an initial low-cost routing, rather than to establish optimality from a high-cost, feasible routing. In addition, the heuristic allows us to run the routing model in real time.

2.4.3 Insertion heuristic to determine opportunity cost

Now that we have developed suitable routings for sample scenarios (X^S_{T+1}, W^S), we want to compute opportunity cost for selling any product $j = (f, z)$ at price set p^f, where $f \in \bar{F}_t$ is any flight that has not "arrived" yet. Since we cannot define product j for any non-scheduled flight until the request for the flight actually comes in, the procedure needs to be carried out in real-time (at least for these flights). To achieve the desired efficiency for online calculation, we apply a simple insertion heuristic that computes the cost $\Delta_j D(X^S_{T+1}, W^S, H)$ of adding product j to the final routings for sample (X^S_{T+1}, W^S). We have

$$\Delta_j D(X^S_{T+1}, W^S, H) := D(X^S_{T+1} \cup (j, p^f), W^S, H) - D(X^S_{T+1}, W^S, H).$$

To compute this cost, we first fix the configuration C' determined in step 1, giving sectors P'. To ensure that we use the most economical feasible configuration, we collapse further sectors in an airspace if this action does not increase k_{acu}, and update C' accordingly. For each product type $z \in Z$ and sample scenario (X^S_{T+1}, W^S), we then determine route $r' \in R^f_z$ that generates the lowest increase in displacement cost while keeping the resulting solution feasible, i.e., the route with lowest $\delta^{f,S}_r$ where

$$\delta^{f,S}_r = \begin{cases} d^f_r & \text{if } f \in F^G \text{ and } \bar{k}_l \leq 1, l \in L' \\ d^f_r & \text{if } f \in F_t \backslash F^G_t \text{ and } \bar{k}_l + w_{frl} \leq 1, l \in L' \\ M & \text{otherwise.} \end{cases}$$

If the resulting traffic flow for product j is infeasible for $r \in R_z^f$, we assign insertion cost of M to the product (which we set to maximum observed displacement cost); else we assign insertion cost of d_r^f. We differentiate between scheduled and non-scheduled flights in our conditions because all scheduled flights are already part of X_{T+1}^S (for all $S \in \mathcal{S}(X)$) so that we do not add w_{frl}. This approach will generate slightly larger opportunity cost for non-scheduled flights, which reflects the additional uncertainty these flights create. The insertion cost $\Delta_j D(X_{T+1}^S, W^S, H) = \delta_{r'}^{S,f}$ is then used in (2.8) to approximate opportunity cost.

The approach described above incorporates the expected effect on total displacement cost $D(X_{T+1}, W, H)$ in the computation of opportunity cost, but it does not reflect the soft considerations discussed in §2.3.2. To ensure *revenue neutrality* and *fairness*, we compute:

$$\Delta_p \epsilon^{RN} = |1 - \hat{p}^f| \tag{2.10}$$

$$\Delta_p \epsilon^{FR} = Var(p^f). \tag{2.11}$$

2.5 Numerical experiments

The objective of this paper is to determine which setting for flight-to-route assignments is most effective in reducing total displacement cost. Therefore, we test our proposed methodology on a medium-sized case study. The quality of the flight-to-route assignments in each setting is measured in terms of a) reduction of displacement cost (both delay and rerouting), b) availability of feasible flight assignments and c) ability to ensure revenue neutrality and fairness. In §2.5.1 we describe the settings and corresponding decision policies which we evaluate, in §2.5.2 we discuss the simulation study and evaluation process, and in §2.5.3–2.5.4 we report our results.

2.5.1 Simulation settings: Decision policies

In order to examine the value of different flight-to-route assignments in the pre-tactical phase, we compare 3 different settings. In the first setting, the NM retains full flexibility on how to route flights through the network, leading to the following decision policy:

- *Network Manager decision (NMd)*: We let the NM choose the route for each flight. This is modeled by assigning the most flexible trajectory product to all flights (at a price of 1 to ensure revenue neutrality), and determining the resulting displacement cost using CILP and the routing heuristic.

In the second setting, the AU retains full flexibility to choose the route for each flight, giving the second decision policy:

- *Airspace User decision (AUd)*: We let the AU choose the route for each flight. This is modeled by assigning the least flexible trajectory product to all flights (at a price of 1 to ensure revenue neutrality), and determining the resulting displacement cost using CILP and the routing heuristic. We implicitly assume that all AU prefer to fly the shortest route, which seems reasonable given that the charge is the same for all routes.

In the third setting, flexible trajectory products are introduced (see §2.3.1) and the NM decides on the price charged to the AU for each product type. Given the price, the AU then decides which product to purchase (based on the AU choice model) and the NM finally decides how to route the flight based on the purchased product type. We differentiate three pricing policies in this setting:

- *Foresight static pricing (FS)*: We set initial prices for each product type and keep them constant across the booking horizon. In order to provide a fair comparison to the dynamic pricing policies, we set prices for each product type to the average price offered to AUs for the respective product under the FD policy (see below). These prices are evaluated through the simulation runs described in the next subsection, which is why the policy implicitly includes foresight. Using the FS policy lets us determine the value of making dynamic decisions.

- *Hindsight dynamic pricing (HD)*: Instead of using static prices, we adjust prices dynamically based on the flight arriving to the booking process. To compute prices, we determine insertion cost based only on existing bookings; hence the term hindsight policy. In particular, we determine the cost of inserting each product type into the best routing of all current flights. Using the HD policy lets us determine the value of simulating the future (by comparing it to the FD policy below), which is computationally expensive. In summary, prices for the HD policy are set based on (2.6), where we compute expected opportunity cost as:

$$\Delta_{(j,p)}\ \hat{V}_{t+1}(X) = \Delta_j D(X_t, W^S, H) + \theta^{RN}\Delta_p \epsilon^{RN} + \theta^{FR}\Delta_p \epsilon^{FR}.$$

- *Foresight dynamic pricing (FD)*: In contrast to HD, insertion cost are determined based on the foresight approach described in §2.4.3, where each scenario reflects a full schedule

for cut-off time $T+1$. We compute opportunity cost as follows:

$$\Delta_{(j,p)}\ \hat{V}_{t+1}(X) = \frac{\sum_{S\in\mathcal{S}(X)} \Delta_j D(X^S_{T+1}, W^S, H)}{|\mathcal{S}(X)|} + \theta^{RN}\Delta_p\epsilon^{RN} + \theta^{FR}\Delta_p\epsilon^{FR}.$$

To determine insertion cost in the two dynamic pricing policies (HD and FD), we use the CILP, the routing heuristic and the insertion heuristic from §2.4. Note that all prices are quoted as relative prices with regards to the benchmark price on a certain origin-destination pair (see §2.3.2), so they automatically reflect the route characteristics of a flight.

2.5.2 Simulation study

Simulation runs and evaluation

We conduct a simulation study in order to evaluate the impact of the different settings. In each run, we simulate a full booking horizon, where the sequence of flights arriving to the booking process (sample path) is fixed across pricing policies to ensure comparability. Each sample path ω defines an ordering of flights in $F^\omega = F^G \cup F^{N,\omega}$ where $F^{N,\omega}$ is a subset of non-scheduled flights F^N and varies from one run to another. For each flight $f \in F^\omega$ arriving to the booking process, we determine the price vector p^f_t based on the respective pricing policy. For all dynamic pricing policies, we compute opportunity cost from a fixed number of scenarios $S \in \mathcal{S}(X)$. Each scenario consists of a set of flights and product types for the end of the booking horizon (at $T+1$) as well as sector capacities based on W. In our simulation, we evaluate a total of 20 scenarios each time to compute opportunity cost, and update the choice of scenarios every 10 flights. These parameters have proven to provide the best trade-off between solution quality and solving time in initial simulation runs. The scenarios vary for each run, but are fixed across policies.

In our simulation study we choose to offer two product types, one flexible and one direct product: $Z = \{flex, dir\}$. When a pricing decision p^f has been made for a booking request (giving p_{flex} and p_{dir}), we model the product decision according to the AU choice model. In particular, we assume that the AU chooses between the two product types depending on the ratio of their prices, $\nu = p_{flex}/p_{dir}$. We use a binary logit function to model AU choice where the probability of choosing one product type over another changes rapidly close to an

infliction point for ν (which we set to 0.85):

$$P_{flex}(p^f) = \frac{\exp^{(30-30\nu/0.85)}}{\exp^{(30-30\nu/0.85)}+1} = 1 - P_{dir}(p^f).$$

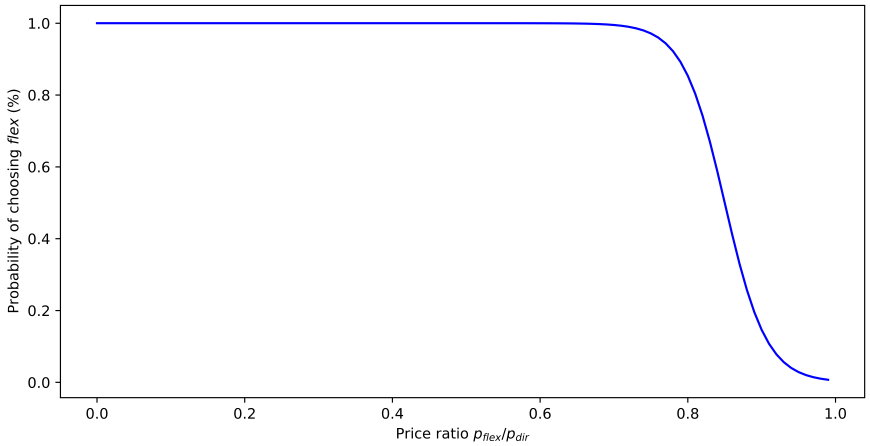

Figure 2.2: Binary logit function to model AU choice

Given all product choices for $t = 1, \ldots, T$ for a given sample path, we then determine displacement cost of the resulting state X_{T+1}, given a further realization of uncertainty W_{T+1}), using the CILP and Algorithm 1. We also obtain whether the routing of X_{T+1} is feasible, and if it is not, how many flights could not be assigned. For the FS policy, we set static unit prices $p_{flex} = 0.98$ and $p_{dir} = 1.16$ based on the simulation results from the FD policy. We simulate a total of 2,000 runs for each policy.

Description of case study

The case study used for the computational analysis is based on Starita et al. (2016). The artificial network is depicted in Figure 2.3. It consists of 5 airspaces (Q, R, S, T, U), four of which have 2 elementary sectors, and one (airspace U) has 3 elementary sectors. The airspaces are arranged such that the shortest routes (i.e., routes b and c in Fig. 2.3) always cross airspace Q. Capacity budget H is 10 sector-hours for airspace U and 7 sector-hours for all others. Furthermore, the network contains a total of 120 scheduled flights and a pool of 80 non-scheduled flights from which we sample the subsets $F^{N,\omega}$. The size of subsets $F^{N,\omega}$, for all sample paths and scenarios $S \in \mathcal{S}(X)$, is drawn from a normal distribution with mean 30 and standard deviation of 8. Each flight has assigned to it one of 3 aircraft types, one of 8 origin-destination pairs and a departure time. Each origin-destination pair is associated with

8-10 different route options which, in turn, are fully described by the sequence of elementary sectors they cross during a certain time after departure, and a delay (capped at 30 minutes). The displacement cost of any flight varies between 0 and 1750 EUR and depends on aircraft type (small, medium or large aircraft) and route choice. It increases with aircraft size, with the length of the route and with delay. To account for infeasible routings, a dummy route is added for each origin-destination pair with no sector requirements and displacement cost twice the largest cost on any non-dummy route. To model the two product types, we define $R^f_{dir} := r^*_f$ (where r^*_f is the shortest route of flight f) and $R^f_{flex} := R^f$.

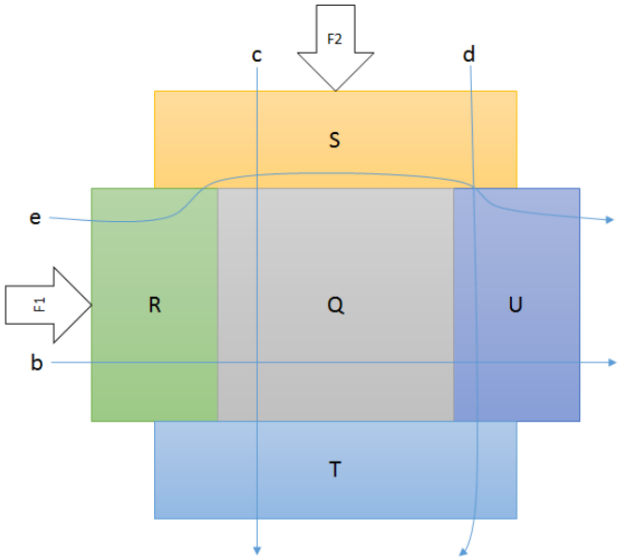

Figure 2.3: Air route network with 2 exemplary flights.
Figure taken from Starita et al. (2016).

We consider an operating time interval of 2 hours, with configuration changes allowed every 30 minutes. To model capacity uncertainty, we reduce the capacity of a randomly chosen elementary sector by 10% in 5% of cases and by 30% in another 5% of cases, for each airspace. This reduction applies across the 2 hour operating time interval. We provide the full case study dataset in the appendix to allow results to be replicated.

2.5.3 Performance of routing heuristic

The simulation study allows us to evaluate how well the different static and dynamic pricing policies are working. Since the performance of all dynamic policies depends on how well we can estimate opportunity cost, we first comment on the performance of the proposed

heuristic (CILP and Algorithm 1) to generate these estimates. All algorithmic approaches are implemented using Python, supported by the commercial Gurobi solver for optimization models. The simulations are run on AWS Batch using 4 GB RAM.

We test the heuristic by computing $D(X^S_{T+1}, W^S, H)$ on a total of 20 scenarios, where X^S_{T+1} and W^S are generated as described in §2.4.3. Note that we assume product type $flex$ for all flights (i.e., routes can be chosen from R^f for all f). The results from the heuristic are compared against the exact ILP solution determined using Gurobi solver, and are summarised in Table 2.1 below. The heuristic determines routings with an average displacement cost gap of 11.3% to the optimum solution based on the ILP over 20 instances. An average of 4.4% and 5.6% of flights have to be assigned to dummy routes for the ILP and heuristic, respectively.

Table 2.1. Simulation results of routing model on 20 instances.

	ILP			*Heur.*			
	Cost	Time (s)	N. a.	Cost	Time (s)	N. a.	Cost gap
1	31,679	1,719	5.3%	35,695	2.6	7.3%	12.7%
2	28,000	3,953	5.3%	30,668	2.3	6.7%	9.5%
3	26,880	3,256	3.3%	30,494	2.2	6.7%	13.4%
4	25,254	3,125	4.7%	27,554	2.2	6.0%	9.1%
5	22,724	3,078	2.7%	25,494	2.1	4.0%	12.2%
6	18,687	614	2.0%	20,733	1.7	2.7%	10.9%
7	35,304	10,700	8.7%	38,514	2.7	8.7%	9.1%
...	...	...	...	...	...	...	...
20	30,737	4,813	4.7%	33,376	2.6	6.0%	8.6%
⌀	26,023	3,022	4.4%	28,852	2.2	5.6%	11.3%

N.a. stands for flights that are not assigned to a non-dummy route.

While the heuristic approximates displacement cost reasonably well when compared to the optimum solution, we are eventually interested in how well it performs at estimating opportunity cost. Figure 2.4 compares the resulting opportunity cost between the ILP and heuristic for all 200 flights in the case study. The cost are computed by inserting the flight into the 20 schedules from Table 2.1 determined through either the ILP or heuristic (see (2.8)). The strong correlation with R^2 of 98% suggests that the heuristic represents an effective method to approximate actual opportunity cost. In particular, we observe a mean absolute error (MAE) of 123, which is less than 10% of average opportunity cost of 1,250.

Furthermore, the approach offers the necessary efficiency for real-time implementation. The heuristic solves all 20 instances significantly faster than the ILP with an average solution time of 2.2 seconds, compared to 3,022 seconds $\approx$ 50 minutes for the exact approach. Recall

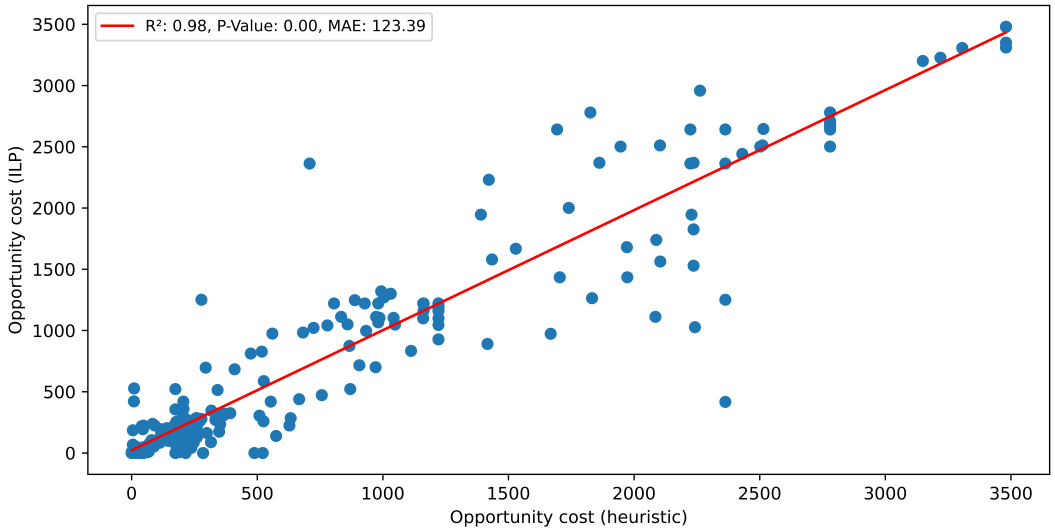

Figure 2.4: Opportunity cost of ILP vs heuristic on 200 flights

that the heuristic has complexity $\mathcal{O}(m(n-g)^2+mn)$ for each scenario, where m is the number of sectors opened across time (approx. 40 in the case study), n the number of flight-route combinations (approx. 1,350) and g the number of flights (approx. 150). With the heuristic, a realistic-sized instance in European airspace with 25,000 daily flights at 7 route options each (i.e., 175,000 flight-route combinations), and around 2,000 sectors opened across time would still require 20 days to solve. However, Melgosa et al. (2019) find that on a typical day more than 75% of flights will fly the shortest route, and are thus unaffected by reroutings. To be conservative, we estimate that 50% of flights can always be assigned to their shortest route without affecting other flights, which reduces the number of flights to include in the optimization. If the day of operation is further split into four roughly equally sized time interval (e.g., morning, afternoon, evening and night), and the routing problem solved for each time interval separately, a full instance (i.e., scenario) could be solved in less than 2 hours. Since each scenario can be evaluated independently, the process can be parallelized across scenarios, which allows computing reliable opportunity cost multiple times a day.

2.5.4 Comparison of settings

Using the routing heuristic to determine opportunity cost, we test how effective the different settings and policies are at reducing total displacement cost. For all policies with flexible products (FS, HD and FD), we set penalties θ^{RN} and θ^{FR} to a value 10 times the maximum possible displacement cost per flight (i.e., 1750), which has shown to provide the best results.

The following table compares displacement cost, relative cost savings, share of flights that are not assigned (in case of infeasibility), and ATM revenues (relative to capacity cost) for the five policies.

Table 2.2. Simulation results of pricing policies over 2,000 runs.

Setting	Policy	Cost	Cost savings of FD	Not assigned	Revenues
Flexible products	HD	54,097	-40%*	21.5%	100.0%
	FS	40,297	-20%*	14.6%	105.9%
	FD	32,248	-	10.2%	100.0%
No choice	NMd	27,475	+17%*	6.6%	100.0%
Full choice	AUd	83,950	-62%*	31.3%	100.0%

* Significant at 95% confidence level.

As expected, the lowest displacement cost of 27,475 are generated if the network manager is given full flexibility to assign flights to routes (NMd). These cost roughly triple when shifting the decision power entirely to the airspace user (AUd), with cost of 83,950. All three policies with flexible trajectory products report substantially lower cost than the AUd (or full-choice) policy, showing that the value of providing such products is high. With cost of 32,248 the foresight dynamic policy (FD) reduces cost vs the AUd even by 62% (significant at 95% level), and even comes close to the NMd performance. Despite the small size of the case study, the results therefore suggest that providing flexible trajectory products can serve as a powerful tool to balance AU choice with cost-efficient routing. Among the three flexible product policies, we can see that the FD significantly outperforms the HD policy (which computes opportunity cost based only on existing bookings), with an average reduction of 40%. This confirms that using a foresight approach provides significant additional value. Furthermore, FD also significantly outperforms the static pricing policy (FS), showing that there is further value in dynamically adjusting prices for flexible products. Note that the FS policy even reports lower cost than the HD policy. This suggests that there is no virtue in merely offering dynamic prices if these prices do not adequately reflect the expected opportunity cost generated by the product.

In the optimum case, an average of 6.6% of flights have to be assgined to dummy routes, reflecting infeasible routings. As expected, the flexible product policies show relatively lower shares on dummy routes (10-22%) than the AUd policy (31%). However, the absolute level of these shares depend largely on the structure of the case study and can therefore not be directly interpreted. Finally, total revenues generated by the both dynamic pricing policies (HD and FD) cover quite exactly 100.0% of capacity cost, which suggests that requirement

(2.10) is in fact very effective at ensuring revenue neutrality. However, in the case of FS the collected revenues exceed capacity cost by a sizable 5.9%. Since prices in the FS are static and hence cannot be adjusted based on collected revenues within the booking period, the challenge is to determine appropriate prices in advance that will create cost-neutral revenues (possibly at the expense of displacement cost). For NMd and AUd policies, the NM will simply set the price of 1 for the chosen route option and thus retains full control to manage revenue neutrality. In the following sections, we discuss managerial insights specifically for the NM and AU.

Managerial insights for Network Manager

The main objective of the NM is to design a route assignment mechanism that leads to lowest possible displacement cost, and can be implemented with reasonable resources. To judge the resource consumption of the different policies, we compare computing times for offline and online calculations. The offline calculation (i.e., determining best routings for all scenarios) is performed by the HD policy in 0.2 seconds, and by the FD policy in 32 seconds. The gap is largely due to the fact that we only need to evaluate one rather small scenario (i.e., all existing bookings) for HD, but 20 full scenarios for FD. The NM therefore faces a trade-off between solution quality (40% lower cost for FD) and computing resources (150x longer computing time vs HD). While the current computation time of 32 seconds for FD does not pose any problems for implementation, it increases roughly quadratically with the size of the case study (§2.5.3). As long as the offline calculation, parallelized across scenarios, can be carried out in 24 hours or less, the FD policy is prefered due to its superior performance. If computing time renders a daily update of the FD policy infeasible, static pricing (FS) is prefered over the HD and FD policies. The online calculation (i.e., determining optimal price set based on insertion cost) is performed in 0.07 and 0.2 seconds for the HD and FD policies, respectively. This renders both policies suitable for real-time application. Most importantly, the online solution time only increases with the number of scenarios, but not with the size of the case study, ensuring that the FD remains sufficiently fast for larger networks.

To analyse how effectively each policy reduces displacement cost, we compare in Figure 2.5 the distribution of total displacement cost for the NMd, FD and AUd policies aross all runs. In case the AU fully decides on the route option (AUd), the displacement cost vary widely across runs. This is because a large share of flights are unable to be routed feasibly according to the chosen trajectory product and are assigned to dummy routes with large penalties. In contrast, the NMd and FD produce similarly narrow distributions, with cost

varying roughly between 25,000 and 40,000 for half of the runs. The results suggest that the NM can expect a similarly stable cost distribution under dynamic trajectory pricing (FD) as when they fully decide on the routing themselves (NMd), which gives confidence in the mechanism.

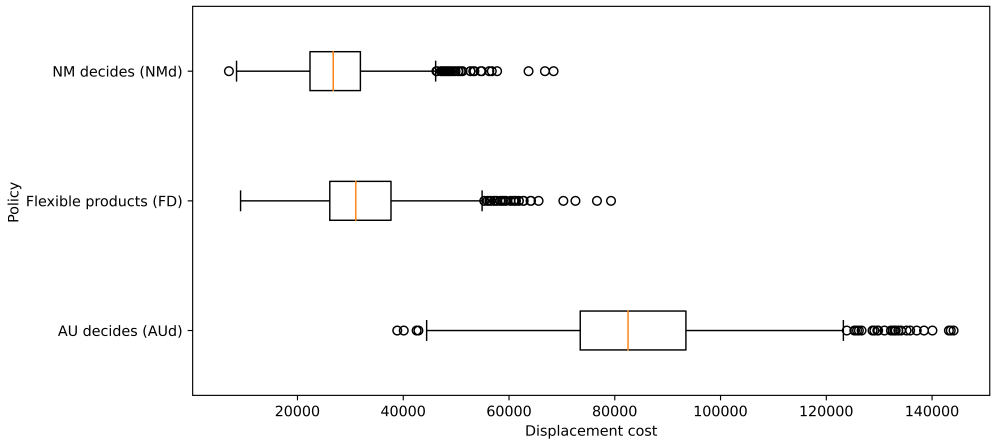

Figure 2.5: Displacement cost distribution across all runs (n = 2,000 runs)

Finally, we compare the average opportunity cost for the two product types to judge the "value of flexibility" of offering a *flexible* product. Under the FD policy, the opportunity cost for the *direct* and *flexible* product amount to 1,475 and 510, respectively. While the absolute numbers depend on the case study design and cannot be interpreted, the magnitude of the difference (almost 3 times lower cost for flexible products) shows that there is significant value in providing flexible products in the pre-tactical ATM process.

Managerial insights for Airspace User

The main objective of the AU is to keep total cost from delays, detours and ATM service charges as low as possible while retaining maximum route choice in the booking process. We therefore analyse how displacement cost are distributed among flights in Figures 2.6 and 2.7. In contrast to Figure 2.5, where we compare total displacement cost between runs, Figure 2.6 compares average displacement cost per flight between all 200 flights in the case study. We an see that under the FD policy, small aircrafts do not only show the highest median displacement cost, but also the largest variation in cost. This may seem counterintuitive because displacements are generally less costly for smaller than for larger aircraft. However, due to their relatively cheaper displacements, these aircraft are more likely to be rerouted or delayed so as to avoid congested sectors. In contrast, the optimization will

avoid displacing larger aircraft altogether due to the high cost it creates. The result suggests that business aviation and non-scheduled flights operating on small aircraft should expect the largest uncertainty regarding delays and detours when planning their flights.

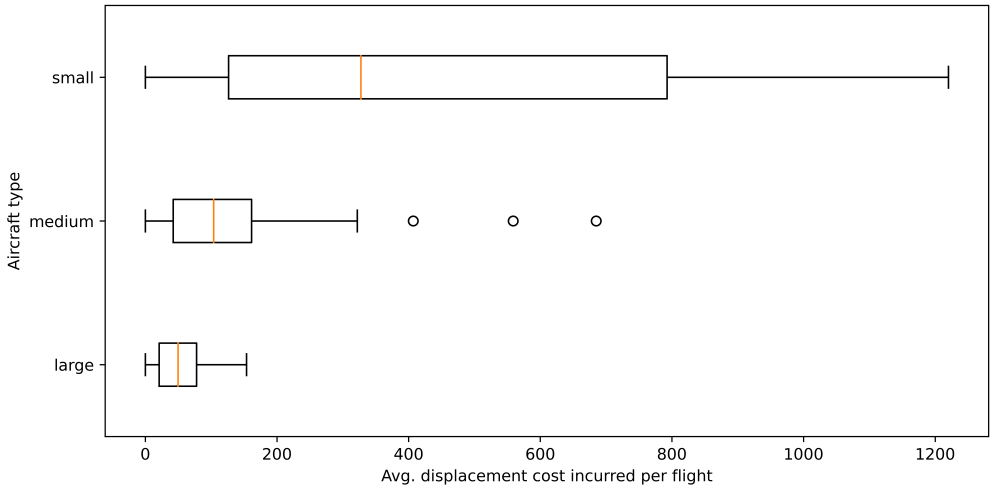

Figure 2.6: Distribution of average displacement cost on DPTF policy by aircraft type

Figure 2.7 compares the distribution of displacement cost across flights for different policies. Similar to the earlier observation, it shows that AUs can expect a much more unequal distribution of displacement cost if they themselves decide on the desired route (AUd). In fact, 75% of flights show displacement cost below 200 for the NMd and FD policies, but reach up to 1,200 under AUd. Generally, the distribution of cost is similarly narrow for the NMd and FD policies. This finding suggests that even if the NM can fully decide on the route option, this does not lead to a more unequal distribution of delays and detours among flights. Recall that since the AU has to purchase one of the products offered in the booking process, we implement a *fairness condition* (see §2.3.2) to prevent the NM from choosing prices such that one product type is effectively "imposed" on the AU. With the fairness condition in place, the average price offered for the *fexible* and *direct* product under the FD policy are 0.98 and 1.16, respectively. Given our AU choice model, this price ratio p_{flex}/p_{dir} of 0.84 gives a probability to purchase the flexible product of around 54%, which confirms that indeed no one product type is imposed on the AUs. If the fairness condition was levied, we would expect a probability close to 100%, because the optimization would force the AU to purchase the flexible product whenever its opportunity cost is strictly lower than that for the direct product. Therefore, the fairness condition creates an effective trade-off between network performance (in terms of displacement cost) and AU choice.

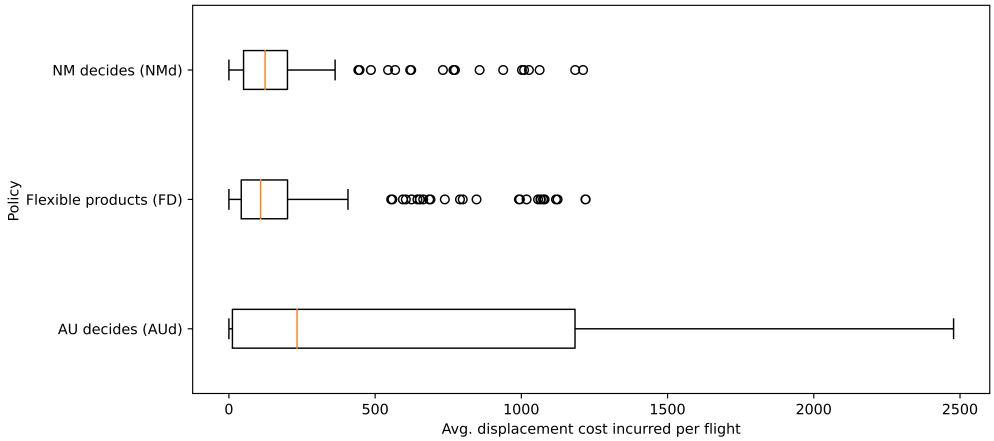

Figure 2.7: Distribution of average displacement cost by policy (n = 200 flights)

Limitations

There are a few limitations to the outlined experiments. In general, the case study represents a rather small artificial network with several design choices that limit the generalizability of results. First, the route options are designed such that the shortest route will almost always go through a bottleneck sector, which explains why the AUd policy performs rather poorly. Second, the displacement cost for each route, albeit founded on recent research, represent artificial values so that their absolute levels cannot be interpreted. Third, the performance of the policies depends on a range of pre-defined case study parameters, such as the capacity budget and the penalty for dummy routes. For instance, a higher budget would lead to less infeasible routes, fewer penalties and thus smaller differences in performance among policies. Apart from the case study, there are a few modeling assumptions that impact our results. First, our choice model assumes that AUs make their product choice based only on ATM charges, which neglects the influence of fuel cost and delays. Second, we assume that all AUs behave according to the same choice model, but different business models (e.g., legacy vs low-cost carrier) may warrant different behaviors. Finally, we differentiate only two product types in the study, but further products may tilt the results in favor of one or another policy.

2.6 Conclusions

We demonstrate the value of dynamic pricing to steer demand in pre-tacical ATM. European ATM suffers from severe demand-capacity imbalances, leading to EUR 550 million in ATM-related delay cost in 2017, according to Eurocontrol (2018). To reduce these imbalances, we

develop a methodological framework to test flexible trajectory products that are offered to AUs in the booking process and differ in how flexibly the NM can route the flights during departure day. Pricing the trajectory products is particulary challenging since it requires solving a hard routing problem in the terminal condition, which itself is subject to uncertainty around arriving fligths and capacity provision. We use a MMKP-based heuristic to solve the difficult routing and sector-opening problem (offline), and an insertion heuristic with foresight to capture the uncertainties (online). By separating the offline generation of routing schedules from the online estimation of opportunity cost, we can compute trajectory prices in real-time even for realistic-sized instances. Once opportunity cost are estimated, they be used as an input to various dynamic pricing policies. The proposed dynamic policies are tested on an artificial network with 150 flights, and compared against two other settings where either the NM or AU retains full mandate to route flights. The results show that dynamic pricing with foresight leads to network performance almost as high as if the NM decides on routing, and significantly outperforms both static pricing and full AU choice. The main advantage of the proposed trajectory pricing scheme in ATM is that it helps steer demand in line with capacity while still providing choice to the AU in the booking process.

Acknowledgements

This project has received funding from the SESAR Joint Undertaking within the framework of SESAR 2020 and the EU's Horizon 2020 research and innovation programme under the Grant Agreement Number 893380. The opinions expressed herein reflect the author's view only. Under no circumstances shall the SESAR Joint Undertaking be responsible for any use that may be made of the information contained herein.

Appendix

2.A Simulation results

Table 2.A.1. Simulation results of routing model on 20 instances.

	ILP			*Heur.*			
	Cost	Time (s)	N. a.	Cost	Time (s)	N. a.	Cost gap
1	31,679	1,719	5.3%	35,695	2.6	7.3%	12.7%
2	28,000	3,953	5.3%	30,668	2.3	6.7%	9.5%
3	26,880	3,256	3.3%	30,494	2.2	6.7%	13.4%
4	25,254	3,125	4.7%	27,554	2.2	6.0%	9.1%
5	22,724	3,078	2.7%	25,494	2.1	4.0%	12.2%
6	18,687	614	2.0%	20,733	1.7	2.7%	10.9%
7	35,304	10,700	8.7%	38,514	2.7	8.7%	9.1%
8	35,247	3,672	7.3%	39,313	2.8	9.3%	11.5%
9	29,772	2,832	6.0%	31,632	2.1	7.3%	6.2%
10	15,388	504	0.7%	17,581	1.7	1.3%	14.3%
11	18,328	739	2.0%	20,885	1.9	2.7%	14.0%
12	26,243	3,591	4.7%	29,970	2.4	6.7%	14.2%
13	21,046	705	2.7%	23,910	1.9	3.3%	13.6%
14	25,525	823	4.0%	27,434	2	5.3%	7.5%
15	29,793	3,553	5.3%	33,078	2.4	7.3%	11.0%
16	23,676	1,178	2.7%	27,041	2.2	4.0%	14.2%
17	16,876	858	2.0%	19,775	1.8	2.0%	17.2%
18	42,084	10,121	12.0%	45,013	2.5	13.3%	7.0%
19	17,220	603	1.3%	18,875	1.7	2.0%	9.6%
20	30,737	4,813	4.7%	33,376	2.6	6.0%	8.6%
⌀	26,023	3,022	4.4%	28,852	2.2	5.6%	11.3%

N.a. stands for flights that are not assigned to a non-dummy route.

2.B Case study data

Table 2.B.1. Overview of routes for case study.

OD	(Elementary sector, time period)						Delay	Displacement cost			z
	1	2	3	4	5	6		small	med.	large	
0	(0, 2)	(5, 2)	(6, 2)	(4, 2)			0	0	0	0	0
0	(0, 2)	(3, 4)	(4, 2)				0	152	280	355	1
0	(0, 2)	(3, 2)	(6, 2)	(4, 2)			0	69	127	162	1
0	(0, 2)	(5, 2)	(3, 2)	(4, 2)			0	69	127	162	1
0	(0, 2)	(5, 2)	(6, 2)	(4, 2)			1	48	90	100	1
0	(0, 2)	(5, 2)	(6, 2)	(4, 2)			2	120	236	313	1
0	(0, 2)	(5, 2)	(6, 2)	(4, 2)			3	204	450	560	1
0	(0, 2)	(5, 2)	(6, 2)	(4, 2)			4	321	693	888	1
0	(0, 2)	(5, 2)	(6, 2)	(4, 2)			5	453	1004	1275	1
0	(0, 2)	(5, 2)	(6, 2)	(4, 2)			6	611	1390	1740	1
1	(1, 2)	(10, 4)	(4, 2)				0	0	0	0	0
1	(1, 2)	(6, 4)	(5, 2)				0	152	280	355	1
1	(1, 2)	(10, 4)	(4, 2)				1	48	90	100	1
1	(1, 2)	(10, 4)	(4, 2)				2	120	236	313	1
1	(1, 2)	(10, 4)	(4, 2)				3	204	450	560	1
1	(1, 2)	(10, 4)	(4, 2)				4	321	693	888	1
1	(1, 2)	(10, 4)	(4, 2)				5	453	1004	1275	1
1	(1, 2)	(10, 4)	(4, 2)				6	611	1390	1740	1
2	(2, 1)	(3, 1)	(8, 2)	(10, 2)	(6, 1)	(7, 1)	0	0	0	0	0
2	(2, 1)	(3, 1)	(0, 2)	(1, 2)	(6, 1)	(7, 1)	0	152	280	355	1
2	(2, 1)	(3, 1)	(8, 2)	(1, 2)	(6, 1)	(7, 1)	0	69	127	162	1
2	(2, 1)	(3, 1)	(8, 2)	(10, 2)	(6, 1)	(7, 1)	1	48	90	100	1
2	(2, 1)	(3, 1)	(8, 2)	(10, 2)	(6, 1)	(7, 1)	2	120	236	313	1
2	(2, 1)	(3, 1)	(8, 2)	(10, 2)	(6, 1)	(7, 1)	3	204	450	560	1
2	(2, 1)	(3, 1)	(8, 2)	(10, 2)	(6, 1)	(7, 1)	4	321	693	888	1
2	(2, 1)	(3, 1)	(8, 2)	(10, 2)	(6, 1)	(7, 1)	5	453	1004	1275	1
2	(2, 1)	(3, 1)	(8, 2)	(10, 2)	(6, 1)	(7, 1)	6	611	1390	1740	1
3	(2, 1)	(3, 1)	(9, 2)	(10, 2)	(6, 1)	(7, 1)	0	0	0	0	0
3	(2, 1)	(3, 1)	(4, 2)	(5, 2)	(6, 1)	(7, 1)	0	152	280	355	1
3	(2, 1)	(3, 1)	(9, 2)	(5, 2)	(6, 1)	(7, 1)	0	69	127	162	1
3	(2, 1)	(3, 1)	(9, 2)	(10, 2)	(6, 1)	(7, 1)	1	48	90	100	1
3	(2, 1)	(3, 1)	(9, 2)	(10, 2)	(6, 1)	(7, 1)	2	120	236	313	1
3	(2, 1)	(3, 1)	(9, 2)	(10, 2)	(6, 1)	(7, 1)	3	204	450	560	1
3	(2, 1)	(3, 1)	(9, 2)	(10, 2)	(6, 1)	(7, 1)	4	321	693	888	1
3	(2, 1)	(3, 1)	(9, 2)	(10, 2)	(6, 1)	(7, 1)	5	453	1004	1275	1

...											
3	(2, 1)	(3, 1)	(9, 2)	(10, 2)	(6, 1)	(7, 1)	6	611	1390	1740	1
4	(4, 2)	(9, 2)	(8, 2)	(0, 2)			0	0	0	0	0
4	(4, 2)	(3, 4)	(0, 2)				0	152	280	355	1
4	(4, 2)	(9, 2)	(3, 2)	(0, 2)			0	69	127	162	1
4	(4, 2)	(3, 2)	(8, 2)	(0, 2)			0	69	127	162	1
4	(4, 2)	(9, 2)	(8, 2)	(0, 2)			1	47	90	100	1
4	(4, 2)	(9, 2)	(8, 2)	(0, 2)			2	120	236	313	1
4	(4, 2)	(9, 2)	(8, 2)	(0, 2)			3	204	450	560	1
4	(4, 2)	(9, 2)	(8, 2)	(0, 2)			4	321	693	888	1
4	(4, 2)	(9, 2)	(8, 2)	(0, 2)			5	453	1004	1275	1
4	(4, 2)	(9, 2)	(8, 2)	(0, 2)			6	611	1390	1740	1
4							0	1221	2780	3480	0
5	(4, 2)	(0, 4)	(1, 2)				0	0	0	0	0
5	(5, 2)	(6, 4)	(1, 2)				0	152	280	355	1
5	(4, 2)	(0, 4)	(1, 2)				1	47	90	100	1
5	(4, 2)	(0, 4)	(1, 2)				2	120	236	313	1
5	(4, 2)	(0, 4)	(1, 2)				3	204	450	560	1
5	(4, 2)	(0, 4)	(1, 2)				4	321	693	888	1
5	(4, 2)	(0, 4)	(1, 2)				5	453	1004	1275	1
5	(4, 2)	(0, 4)	(1, 2)				6	611	1390	1740	1
6	(7, 1)	(6, 1)	(10, 2)	(8, 2)	(3, 1)	(2, 1)	0	0	0	0	0
6	(7, 1)	(6, 1)	(1, 2)	(0, 2)	(3, 1)	(2, 1)	0	152	280	355	1
6	(7, 1)	(6, 1)	(1, 2)	(8, 2)	(3, 1)	(2, 1)	0	69	127	162	1
6	(7, 1)	(6, 1)	(10, 2)	(8, 2)	(3, 1)	(2, 1)	1	47	90	100	1
6	(7, 1)	(6, 1)	(10, 2)	(8, 2)	(3, 1)	(2, 1)	2	120	236	313	1
6	(7, 1)	(6, 1)	(10, 2)	(8, 2)	(3, 1)	(2, 1)	3	204	450	560	1
6	(7, 1)	(6, 1)	(10, 2)	(8, 2)	(3, 1)	(2, 1)	4	321	693	888	1
6	(7, 1)	(6, 1)	(10, 2)	(8, 2)	(3, 1)	(2, 1)	5	453	1004	1275	1
6	(7, 1)	(6, 1)	(10, 2)	(8, 2)	(3, 1)	(2, 1)	6	611	1390	1740	1
7	(7, 1)	(6, 1)	(10, 2)	(9, 2)	(3, 1)	(2, 1)	0	0	0	0	0
7	(7, 1)	(6, 1)	(5, 2)	(4, 2)	(3, 1)	(2, 1)	0	152	280	355	1
7	(7, 1)	(6, 1)	(5, 2)	(9, 2)	(3, 1)	(2, 1)	0	69	127	162	1
7	(7, 1)	(6, 1)	(10, 2)	(9, 2)	(3, 1)	(2, 1)	1	47	90	100	1
7	(7, 1)	(6, 1)	(10, 2)	(9, 2)	(3, 1)	(2, 1)	2	120	236	313	1
7	(7, 1)	(6, 1)	(10, 2)	(9, 2)	(3, 1)	(2, 1)	3	204	450	560	1
7	(7, 1)	(6, 1)	(10, 2)	(9, 2)	(3, 1)	(2, 1)	4	321	693	888	1
7	(7, 1)	(6, 1)	(10, 2)	(9, 2)	(3, 1)	(2, 1)	5	453	1004	1275	1
7	(7, 1)	(6, 1)	(10, 2)	(9, 2)	(3, 1)	(2, 1)	6	611	1390	1740	1

Column name "z" represents product types (0 = *direct*, 1 = *flexible*).

Table 2.B.2. Overview of flights for case study.

Flight ID	OD	Aircraft type	Dep. time	Flight type
F0	0	0	0	S
F1	0	1	1	U
F2	0	0	2	S
F3	0	0	2	S
F4	0	1	3	S
F5	0	1	3	S
F6	0	2	4	S
F7	0	2	4	U
F8	0	0	4	S
F9	0	0	5	S
F10	0	0	5	U
F11	0	1	5	S
F12	0	1	6	S
F13	0	1	7	S
F14	0	1	7	S
F15	0	0	8	U
F16	0	0	9	S
F17	0	0	10	S
F18	0	1	10	S
F19	0	1	12	S
F20	0	1	12	S
F21	0	1	13	S
F22	0	1	14	S
F23	0	1	15	U
F24	0	0	16	S
F25	0	0	17	S
F26	0	0	17	S
...				
F190	2	2	5	U
F191	6	0	23	U
F192	3	2	12	U
F193	4	1	0	U
F194	5	1	3	U
F195	0	2	13	U
F196	4	2	2	U
F197	0	2	2	U
F198	3	2	8	U
F199	5	1	1	U

Flight type S represents scheduled and U non-scheduled flights.

Table 2.B.3. Overview of configurations for case study.

Sector ID	Airspace	Conf. ID	Elementary sector			Capacity
			1	2	3	
P0	R	C1	0	1		18
P1	R	C2	0			18
P2	R	C2	1			18
P3	S	C3	2	3		19
P4	S	C4	2			17
P5	S	C4	3			17
P6	T	C5	4	5		18
P7	T	C6	4			18
P8	T	C6	5			18
P9	U	C7	6	7		19
P10	U	C8	6			17
P11	U	C8	7			17
P12	Q	C9	8	9	10	17
P13	Q	C10	8			16
P14	Q	C10	9	10		18
P15	Q	C1	8	9		17
P16	Q	C1	10			17
P17	Q	C12	8	10		18
P18	Q	C12	9			16
P19	Q	C13	8			16
P20	Q	C13	9			16
P21	Q	C13	10			17

Sector ID refers to collapsed (not elementary) sectors.

Chapter 3

Capacity management: Cross-border capacity planning under uncertainty

[1]

[1]Reviewed and resubmitted in June 2022 to *Transportation Science* under "Cross-border capacity planning in air traffic management under uncertainty" with co-authors Arne K. Strauss, Nikola Ivanov, Radosav Jovanovic, Frank Fichert and Stefano Starita

Abstract

In European air traffic management (ATM), it is an important decision how much capacity to provide for each airspace, and it has to be made weeks or even months in advance of the day of operation. Given the uncertainty in demand that may materialize until then along with variability in capacity provision (e.g., due to weather), Airspace Users could face high costs of displacements (i.e., delays and re-routings) if capacity is not provided where and when needed. We propose a new capacity sharing scheme in which some proportion of overall capacities can be flexibly deployed in any of the airspaces of the same alliance (at an increased unit cost). This allows us to hedge against the risk of capacity underprovision. Given this scheme, we seek to determine the optimum budget for capacities provided both locally and in cross-border sharing that results in the lowest expected network costs (i.e., capacity and displacement costs).

To determine optimum capacity levels, we need to solve a two-stage newsvendor problem: We first decide on capacities to be provided for each airspace, and after uncertain demand and capacity provision disruptions have materialized, we need to decide on the routings of flights (including delays) as well as the sector opening scheme of each airspace to minimize costs. We propose a simulation optimization approach for searching the most cost-efficient capacity levels (in the first stage), and a heuristic to solve the routing and sector opening problem (in the second stage), which is $\mathcal{NP}$-hard. We test our approach in a large-sized simulation study based on real data covering around 3,000 flights across Western European airspace. We find that our stochastic approach significantly reduces network costs against a deterministic benchmark while using less computational resources. Experiments on different setups for capacity sharing show that total costs can be reduced by over 8% if capacity is shared across borders - even though we require that no airspace can operate lower capacities under capacity sharing than without (this is to avoid substitution of expensive air traffic controllers with those in countries with a lower wage level). We also find that the use of a common technology provider is a major obstacle to reap the benefits from capacity sharing, and that sharing capacities across airspaces of the same country may instead be preferred.

Keywords: air traffic management; capacity planning; simulation optimization

3.1 Introduction

We study a two-stage newsvendor problem in the context of air traffic management with cross-border capacity sharing. In the first stage, a network manager needs to decide on capacity levels to be delivered by each capacity provider for their respective airspaces locally, as well as on a (more expensive) cross-border capacity that can be flexibly deployed in multiple airspaces. Cross-border capacity sharing is intended to address underprovisions caused by variations in demand (e.g., through major traffic shifts in the network) and capacity provision (e.g., through weather events or unplanned air traffic controller shortages). In the second stage, flight intentions by aircraft operators materialize and various sources of uncertainty are being resolved. The network manager needs to decide on demand management measures (delay or re-routing) and on the exact sector opening scheme for each airspace to minimize overall costs.

This problem is motivated by various on-going endeavors that advocate a stronger role of the network manager (NM) as well as capacity sharing between air navigation service providers (ANSPs). At present, European ANSPs rely on their individual resources to deliver required capacity levels to meet demand on a day of operation. These resources mostly refer to the number of available air traffic controllers (ATCOs), and since their rosters and shifts are planned several weeks in advance, there are usually only limited options to call in more ATCOs on short notice when additional capacity is needed. In fact, air traffic control capacity and staffing was cited as one of the main reasons for delays in European airspace in 2018 and 2019 (Eurocontrol (2019a, 2020a)). More specifically, about 25% of the 17 million minutes of en-route air traffic flow management (ATFM) delays—i.e., delays imposed by the NM—were attributed to staffing in 2019, with an indication that the lack of ATCOs could have an even higher impact than reported (Eurocontrol, 2020b).

One of the proposed remedies to recurring staffing issues is resource sharing between ANSPs (related to the "capacity-on-demand" concept), which aims to increase resilience to disruptions by enabling a more dynamic temporary delegation of the provision of air traffic services to an alternate provider with spare capacity. The concept forms an integral part of the proposal for the future architecture of the European airspace (SESAR Joint Undertaking, 2019). The Wise Persons Group (2019) also recommends the development of "capacity-on-demand" services, emphasizing the need for improved flexibility and resilience in capacity provision. In fact, cross-border capacity sharing is already at the early implementation stage in Europe, with ANSPs from Finland and Estonia forming an alliance called FINEST. Furthermore, a SESAR project led by ENAIRE (the Spanish ANSP) has tested the sharing

of resources among the Madrid and Sevilla area control centers (ACCs), and between the Palma and Barcelona ACCs; they find that the delegation of ATM services is operationally feasible for low and medium traffic periods (SESAR Joint Undertaking, 2022).

More recently, three European ATM doyens welcome the regulation proposal which strengthens the role of the NM in European ATM by means of, inter alia, "putting the NM in a position to manage the capacity brokering process, including the possibility to facilitate delegation of airspace" (Andribet et al., 2022). Moreover, they suggest that the NM role should be further empowered with four key roles, including the role of capacity manager, which would, "based on pan-European demand-capacity balance analysis, decide on the best measures for a better balance including mandatory delegation of airspace from congested ANSPs to less congested neighboring ANSPs" (Andribet et al., 2022).

Therefore, we focus on a strengthened role of the NM to act as a central decision maker and investigate different settings of capacity sharing in order to gain insights into what might be promising avenues for future European ATM. It is important to emphasize that capacity sharing is not intended to reduce the cost of capacity provision (which differs dramatically between countries in Europe); instead, it is intended to reduce the cost of re-routings and delays stemming from underprovision due to unforeseen events. Clearly, a proposal of substituting ATCOs in a high-wage country with controllers from a low-wage country will likely not be politically acceptable.

Our main contributions are three-fold: (1) we develop a stochastic solution approach to the hard newsvendor problem underpinning capacity planning based on a simulation optimization framework from Andradóttir and Prudius (2010) and a routing heuristic by Künnen and Strauss (2022); (2) we formulate the problem setting with capacity sharing and adjust the solution approach accordingly; and (3) we inform decision makers on how to implement capacity sharing based on results from a large case study using real flight and network data. We find that our stochastic approach significantly improves, as one would expect, on a benchmark that optimizes capacities based on deterministic problem instances. In contrast to existing approaches, our methodology is also suitable to assess the benefits of capacity sharing: since we use capacity sharing to hedge against the risk of capacity underprovision (depending on materialized traffic and weather), its value can only be assessed stochastically across scenarios, and not deterministically for each scenario individually. Despite the strong self-imposed constraint that every airspace has to have at least as much capacity under capacity sharing than without (to make the concept politically more acceptable), we still find that overall savings of over 8% are realistic.

The paper is organized as follows: We review the relevant literature in §3.2 and provide a formal description of the problem in §3.3. In §3.4, we present a stochastic optimization approach to solve the problem efficiently and discuss a method to evaluate capacity sharing. In §3.5, we compare the stochastic approach against a deterministic benchmark on a realistically-sized case study. We also provide different design options for capacity sharing and test them using the proposed methodology. We close with recommendations in §3.6.

3.2 Literature Review

Capacity sharing

The idea of capacity resource sharing is not entirely new, but to the best of the authors' knowledge, there are no academic papers which explore the potential benefits of such a concept in ATM. While Ivanov et al. (2019) evaluate European ATM with a strengthened role for the network manager, they only argue that capacity sharing could improve the cost-efficiency of the system, and do not explore this option further. The report of the European Commission (2020) distinguishes three implementation settings for capacity sharing as part of the future European airspace architecture, namely: non-tactical, time-critical and virtual center settings. Some arrangements in Europe already show certain features of these settings. Examples include the Finnish-Estonian dynamic cross-border sectorisation (FINEST), the planned provision of ATM services by the Maastricht Upper Area Control Center for SloveniaControl, and the operation of a virtual center serving both Zurich and Geneva ACCs (European Commission, 2020). In our work, we investigate capacity sharing designs aligned with both the time-critical and virtual center settings and compare it with the setting without capacity sharing.

While we are not aware of any academic publications on capacity sharing in ATM, there is work in electricity markets that analyzes cross-border capacity balancing, see Baldursson et al. (2018). They consider three settings concerning sharing agreements: either no sharing at all, exchange of electricity across borders and lastly a common reserve that multiple countries can draw on (similar to our considered designs for capacity sharing). They find that a focus on overall social welfare (rather than the individual players' costs) is important to incentivize collaboration. This aligns with our choice of assuming a central planner seeking to minimize total costs whilst ensuring that no provider operates less capacity than without capacity sharing.

Newsvendor problem

Structurally, the planning problem that we consider is known as newsboy or newsvendor problem in the inventory management literature. The basic version of this problem is to decide on what quantity of products to order so as to sell them later at an a priori unknown profit (since this will depend on materialization of uncertain demand). Various variations have been studied over the past decades, see Khouja (1999) and Qin et al. (2011) for reviews. In our context, we assume that demand for air traffic services is independent of our capacity decision such that we do not need to model interactions between the two. This seems justifiable since, in practice, demand is largely driven by flight schedules which in turn depend on end customer demand and airport slot availability (but not on en-route airspace capacity). We also make the simplifying assumption that the unit price of capacity, which we measure in sector-hours, is constant and only differentiated based on whether it can be deployed only locally or for capacity sharing. Another feature of newsvendor problems is how the risk attitude of the decision maker is modeled. We settle on the classic assumption of risk-neutrality (meaning that we optimize expected costs), because we assume that the capacity decisions are modeled for a short period only (e.g., a single day) and would be taken frequently.

Sample Average Approximation

The main challenge in solving our problem is that the evaluation of the expectation is expensive. For such problems, sample average approximation (SAA) had been proposed by Kleywegt et al. (2002), who also show that the optimal solution based on SAA converges to the true optimal value with probability 1. The idea is essentially to replace the expectation with a finite sample average – that then represents a deterministic function – and that accordingly can be minimized using deterministic optimization methods. In particular, SAA approaches usually require estimates of the gradient of the objective function so as to employ gradient-based numerical optimization techniques. For an introduction to SAA approaches, see Kim et al. (2015). In our application, the objective is discontinuous and sub-gradients would be computationally very expensive so that we need a derivative-free approach. Theoretical properties (in particular, bounds on SAA's accuracy) of a data-driven newsvendor were more recently studied by Levi et al. (2015). We adopt a random search method called the 'asymptotically optimal set (AOS) framework' of Hu and Andradóttir (2019) that has been designed specifically for minimization of an expectation over a discrete or continuous domain that cannot be evaluated exactly. This approach does not require gradients and has the attractive theoretical feature that, as the name suggests, it ensures asymptotic optimality

in that the best point in the candidate set converges almost surely to the global optimum, and other sub-optimal candidates will ultimately be eliminated from the candidate set with probability 1. It improves on the 'adaptive search with resampling' approach of Andradóttir and Prudius (2010) by including a method of discarding inferior points from the pool of candidates; this is particularly important in our application since the computational effort of re-evaluating a solution is significant.

Deterministic approach

In a model closest to our work, Starita et al. (2020) tackle the same strategic capacity planning model for European ATM. They propose a heuristic approach to determine capacity levels for a given demand and capacity scenario. The authors also develop different policies that define how the optimal capacity levels for each scenario (also referred to as "capacity budgets") can be combined into one capacity decision. In contrast to their approach - which starts with a number of known scenarios of the future and results in a suggested capacity budget - we start from the capacity budget and work towards better budgets by assessing the quality of each budget on random scenarios. In this process, we focus most computational effort on the most promising solutions, while allowing new budget decisions to be discovered and investigated. Approaching the problem as a stochastic problem is particularly important since we want to study capacity sharing across airspaces as a means to hedge against the risk of capacity underprovision. Using the approach of Starita et al. (2020) would mean that we start with a deterministic flight scenario with all uncertainties resolved; in that case, there is no need anymore for (the more expensive) cross-border control since we can simply adjust the local capacities accordingly. As we demonstrate in the numerical results in §2.5, this deterministic approach leads to much worse decisions. To the best of our knowledge, no other research has yet been carried out to study the strategic capacity planning problem.

3.3 Problem Statement

In this section, we define the two-stage newsvendor problem under consideration. We first discuss the domain of the problem in the first stage, both with and without capacity sharing. We then show that the evaluation of the objective at the second stage is both noisy and expensive and thus set the stage for the proposed simulation optimization in the next section.

3.3.1 First stage: Search space and conditions for capacity sharing

We consider the problem $\min_{x\in\mathcal{X}} f(x)$, where $f(x) = \mathbb{E}_S(G(x,S)) + c^T x$; x is a vector of capacity allocations to airspaces (measured in sector-hours) out of the finite set of designs $\mathcal{X}$; S represents a random scenario for the materialization of demand and potential capacity provision disturbances; $G(x,S)$ is the displacement cost function given scenario S and capacity budgets x; and the capacity unit costs are given in vector c. Displacement costs represent the cost of re-routing and/or delaying flights. The exact evaluation of $f(x)$ is impossible due to the expectation over a complex distribution of scenarios, combined with the fact that even a single objective function observation $G(x,S) + c^T x$ is expensive, since $G(x,S)$ is a large integer program. Moreover, $G(x,S)$ is not continuous in x, which means that we are not able to construct derivatives at all feasible solutions.

The domain $\mathcal{X}$ of potential solutions can be considered to be finite because for each unit time interval (e.g., 30 minutes), each airspace has only a finite number of potential configurations that it can operate, and each configuration corresponds to a fixed number of sector-hours. In other words, one can enumerate all combinations of configurations that an airspace may operate over the course of a day, and correspondingly would know the number of sector-hours required for each such combination, meaning that the optimal solution must be in this finite set. Of course, the cardinality of $\mathcal{X}$ is very high and therefore it is computationally intractable to explore the entire domain.

However, we do not need to explore the entire domain since, at the strategic planning stage, we already have information on flights from scheduled carriers; uncertainty in the spatio-temporal distribution of demand mainly stems from non-scheduled flights and capacity provision disruption (due to adverse weather or ATCO shortages). Non-scheduled flights usually amount to no more than 20% of total traffic, so the majority of traffic is known in advance and the overall pattern in terms of likely congested airspaces is known. This allows us to define a sensible search space $\mathcal{X}$ within certain maximum and minimum bounds of sector-hours along each dimension. In practice, the range of sector-hours to consider for each airspace can be reliably reduced to around 50, giving a practical size of the search space of $50^{|A|}$ (where $|A|$ is the total number of airspaces). Furthermore, with 80% of traffic known in advance, we find that the number of scenarios required to establish sensible estimates for $f(x)$ is around 300, in a practical instance with 3,000 flights.

In the following, we assume that this process has been completed to identify domain $\mathcal{X}$. We then need to identify the best solution $x^* \in \mathcal{X}$ with a limited computational budget that we measure in terms of the maximal number of objective function evaluations.

We need some additional notation in order to fully specify the objective function. First, a solution $x := [(x_a)_{a\in A}; (x^0_a)_{a\in A}]$ consists of sector-hour budgets x_a for each airspace a that can only be used for that airspace, as well as sector-hour budgets x^0_a that can be used in any airspace within the alliance that a is part of. Let alliance $g \in G$ consist of airspaces $a \in A_g \subseteq A$, where G denotes the index set of all such alliances. We allow $A_g = A$ since capacities may theoretically be shared across the entire network. With x_a and x^0_a respectively, we distinguish between the physical allocation of capacity within the headquarters of an ANSP in an airspace (meaning where the controller would actually be based), and where their workforce is being virtually deployed. This difference is important since we assume, perhaps conservatively, that cross-border capacity sharing would only be politically acceptable if no ANSP would have labor displaced, i.e. if each airspace a has at least as many sector-hours physically assigned as they would have if there was no capacity sharing at all. In particular, we require that each airspace's total capacity budget must be no less than the threshold capacity $\bar{x}_a$ (which is calculated by a previous model run without capacity sharing and serves as a static parameter in the model):

$$x_a + x^0_a \geq \bar{x}_a \qquad \forall\, a \in A. \tag{3.1}$$

Finally, since the introduction of capacity sharing capability will lead to additional training and license requirements of ATCOs, those controllers eligible to work in other airspaces will need to be compensated at a higher rate. Therefore, we require $c^0_a > c_a$, where c^0_a denotes the cost per sector-hour of providing capacity in any airspace $a' \neq a$ by an ATCO in airspace a. Since the majority of cost associated with capacity sharing are fixed cost (e.g., from training ATCOs), they occur irrespective of whether the ATCO actually works in a non-local airspace a'; the same applies to c^0.

3.3.2 Second stage: Evaluation of displacement costs

In order to evaluate any capacity decision x, we need to determine in the second stage the displacement cost $G(x, S)$ across a range of scenarios S (which reflect the uncertainty in traffic and capacity provision). Recall that displacement costs refer to the cost of re-routing and/or delaying flights. For a given capacity x and scenario S, function $G(x, S)$ represents the minimum displacement costs generated by routing all flights in scenario S (which we denote by F^S) across the air traffic network constrained by capacities x. Therefore, in order to determine $G(x, S)$ we need to jointly decide on the allocation of shared capacity to each

airspace, the sector opening scheme and the routing of flights.

To model the allocation of shared capacities to individual airspaces, we denote by h_a^0 the number of sector-hours that will be virtually deployed in airspace a, and require

$$\sum_{a \in A_g} h_a^0 \leq \sum_{a \in A_g} x_a^0 \qquad \forall\, g \in G.$$

To illustrate this point, $x_a^0 = 10$ would mean that airspace a plans for 10 sector-hours of capacity sharing at their location which can then be flexibly used in any airspace $a' \in A_g$ within the alliance of which airspace a is a member.

Before we present the complete mathematical model on sector opening and routing, we introduce the necessary notation. The sector opening scheme specifies which airspace operates under which capacity setup at each time. Let each airspace a consist of non-overlapping elementary sectors $e \in E^a$. Depending on traffic, these elementary sectors can be combined into larger, collapsed sectors l in pre-defined ways. The way in which sectors can be combined is governed by configurations. Any configuration $c \in C^a$ represents a partitioning of airspace a into collapsed sectors $l \in L^c$; the elementary sectors contained in any sector l are in turn stored in E^l. Since each elementary or collapsed sector is individually controlled by ATCOs, more collapsed sectors require more ATCOs and thus more sector-hours. Let $\bar{h}_{ac}$ denote the number of sector-hours required to run configuration c in airspace a for one time unit. Furthermore, let the day of operations be divided into equally-sized, discrete operating time intervals $u \in U$ (we choose an interval of 30 minutes as this is the usual time required to change a configuration). We then have $\bar{h}_{ac} = |L^c|/2$ since $|L^c|$ is the number of sectors under c and each time unit is half an hour long. Finally, we denote by κ_l^S the capacity of each sector l under scenario S, which corresponds to the maximum flights that can cross the sector in a single time interval. Note that the sector capacities depend on scenario S because uncertainty in capacity provision (e.g., due to weather) can reduce nominal capacity levels.

To model the routing of flights, we denote by R^f the set of all possible routes for flight f, which always includes the shortest trajectory of the flight alongside various delay and rerouting options. We define each route $r \in R^f$ as a sequence of elementary sector- and time-combinations. For this purpose, we denote by $(b)_{f,r,e,u}$ the sector-route incidence matrix, where b_{freu} is 1 if flight f on route r uses elementary sector e at time u, and 0 otherwise. For each route $r \in R^f$, let d_r^f represent the displacement costs generated by routing flight f through r, which are determined net of the shortest trajectory at no delay.

We can then define the displacement cost function $G(x, S)$ as follows:

$$
\begin{aligned}
G(x,S) = \min_{y,z} \; & \sum_{f \in F^S} \sum_{r \in R^f} d_r^f y_r^f \\
\text{s.t.} \; & \sum_{a \in A_g} h_a^0 \leq \sum_{a \in A_g} x_a^0 && g \in G && (3.2)\\
& \sum_{u \in U} \sum_{c \in C^a} \bar{h}_{ac} z_{acu} \leq x_a + h_a^0 && a \in A. && (3.3)\\
& \sum_{f \in F^S} \sum_{r \in R^f} \sum_{e \in E^l} b_{freu} y_r^f z_{acu} \leq \kappa_l^S && a \in A, c \in C^a, l \in L^c, u \in U && (3.4)\\
& \sum_{c \in C^a} z_{acu} = 1 && a \in A, u \in U && (3.5)\\
& \sum_{r \in R^f} y_r^f = 1 && f \in F && (3.6)\\
& h_a^0 \in \mathbb{N}^+ && a \in A\\
& y_r^f \in \{0,1\} && f \in F, r \in R^f\\
& z_{acu} \in \{0,1\} && a \in A, c \in C^a, u \in U.
\end{aligned}
$$

The objective function minimizes total displacement costs for all the flights in scenario S. Decision variable z_{acu} assigns configuration c for airspace a at time u, and variable y_r^f assigns flight f to route r. Constraints (3.2) ensure that only the available shared capacity is used by the airspaces in each alliance, and constraints (3.3) ensure that the configurations used in each airspace do not exceed the airspace capacity budget (including virtually deployed capacity h_a^0). Constraints (3.4) enforce the sector capacities: if we operate under configuration c at u, we restrict the capacity of any sector l in L^c to κ_l. Constraints (3.5) ensure that each airspace operates under one configuration at any time, and constraints (3.6) ensure that exactly one route is assigned to each flight.

If the evaluation of $G(x, S)$ were easy, we could simply evaluate each solution x for a large sample of scenarios and then substitute the expectation in the newsvendor problem with the sample average. However, Künnen and Strauss (2022) show that this problem is $\mathcal{NP}$-hard such that fully exploring the whole solution space is computationally not practical. In fact, even small instances of the problem show prohibitively long run times. Therefore, we propose a simulation optimization approach to decide which solutions to evaluate how often, given a fixed computational budget.

3.4 Simulation Optimization Approach

Our goal is to find a capacity budget x^* that we confidently believe to lead to low expected network costs across random scenarios $S \in \mathcal{S}$. In the following, we denote by $f(x, S)$ the network costs generated by budget x in scenario S, i.e., $f(x, S) = G(x, S) + c^T x$. Here, any scenario $S = (F^S, W^S)$ constitutes a materialization of traffic F^S due to unknown demand from non-scheduled flights, and capacity W^S due to uncertain provision of capacity services. Note that W^S impacts both the actual available sector-hours x^S (due to ATCO shortages) and actual sector capacities κ^S (due to weather). To find a high-quality capacity decision x^* as outlined above, we use a framework that balances the exploration of new budgets x with the exploitation of existing budgets (meaning to measure $f(x, S)$ for promising candidates x on additional scenarios S to increase our confidence in the average performance); see §3.4.1. We propose an efficient method to evaluate $f(x, S)$ in each measurement step, see §3.4.2.

3.4.1 Framework

We base our approach on the Asymptotically Optimal Set framework by Hu and Andradóttir (2019), which we adjust to our problem. The procedure seeks to determine an optimal capacity decision x^* given solution space $\mathcal{X}$, and is summarized in Algorithm 2. The idea is that we iteratively develop a pool $\mathcal{X}_j^* \subset \mathcal{X}$ of promising solutions (where j is the number of sampled solutions x), which is reduced to decision x^* over time. In each iteration i we either evaluate a new candidate $x \in \mathcal{X} \setminus \mathcal{X}_j^*$, or re-assess an existing candidate from this pool on a new scenario $S \in \mathcal{S}$. To trade off exploration with exploitation, we need to decide how frequently we sample new solutions. Over time we want to explore new solutions less frequently to ensure that more effort can be spent on evaluating the quality of encountered promising solutions. As suggested by Hu and Andradóttir (2019), we sample new solutions at iterations $i = \lfloor j^{1.5} \rfloor$, and re-sample existing solutions otherwise.

At iterations i where $i \neq \lfloor j^{1.5} \rfloor$, we sample a new candidate x according to a pre-defined *sampling strategy*. To ensure that we cover the entire solution space, we use Latin Hypercube Sampling, which is described in the Appendix. We then need to decide whether to accept x into the pool of promising candidates, or whether to discard it. For this purpose, we define a benchmark scenario S_1 without disturbances, i.e., we set $|F^{S_1}|$ to the average number of expected flights, and W^{S_1} such that $x^{S_1} = x$ and $\kappa^{S_1} = \kappa$. We then accept a solution x into the pool if the objective function value $f(x, S_1)$ is at most λ worse than our current best candidate, i.e., $f(x, S_1) - \hat{f}_i(x_{i-1}^*) < \lambda$, where $\hat{f}_i(x)$ denotes the average objective value of x

Algorithm 2 Exploration-exploitation framework to seek optimum capacity

Input: Scenarios $\mathcal{S}$, solution space $\mathcal{X}$, sampling and re-sampling strategy
Initialize: $i = 0$, $j = 1$, $\mathcal{X}_0^* = \emptyset$, $\hat{f}_0(x^*) = M$ (large number)
while $\exists x \in \mathcal{X}_j^* : N(x) < N^{max}$ **do**
  $i = i + 1$
  if $i = \lfloor j^{1.5} \rfloor$ **then**
    Select $x \in \mathcal{X}$ based on *sampling strategy*, evaluate $f(x, S_1)$ and set $j = j + 1$
    if $f(x, S_1) - \hat{f}_i(x^*) < \lambda$ **then**
      $\mathcal{X}_j^* = \mathcal{X}_{j-1}^* \cup \{x\}$, $\hat{f}_i(x) = f(x, S_1), N(x) = 1$
    else
      $\mathcal{X}_j^* = \mathcal{X}_{j-1}^*$
    end if
    for $x \in \mathcal{X}_j^*$ (ensure minimum amount of re-sampling) **do**
      if $N(x) < \lceil j^{0.5} \rceil$ **then**
        Set $n = \lceil j^{0.5} \rceil - N(x)$ and evaluate $f(x, S_{N(x)+1}), \ldots, f(x, S_{N(x)+n})$
        Set $\hat{f}_i(x) = \frac{\hat{f}_{i-1}(x)N(x) + \sum_{k=1}^{n} f(x, S_{N(x)+k})}{N(x)+n}$, and $N(x) = N(x) + 1$
      end if
    end for
    for $x \in \mathcal{X}_j^*$ (discard poor solutions) **do**
      if $\hat{f}_i(x) - \hat{f}_i(x^*) > \lambda / j^{0.15}$ **then**
        $\mathcal{X}_j^* = \mathcal{X}_j^* \setminus \{x\}$
      end if
    end for
  else
    Select $x \in \mathcal{X}_j^*$ based on *re-sampling strategy* and evaluate $f(x, S_{N(x)+1})$
    Set $\hat{f}_i(x) = \frac{\hat{f}_{i-1}(x)N(x) + f(x, S_{N(x)+1})}{N(x)+1}$, and $N(x) = N(x) + 1$
    Update $x_i^* = \arg\min_{x \in \mathcal{X}_j^*} \hat{f}_i(x)$ and
  end if
end while
Output: Capacity budget x_i^* with network cost $\hat{f}_i(x^*)$

at iteration i. Furthermore, if a sufficient amount of objective function evaluations have been conducted, we discard existing solutions based on a rejection threshold. If j candidates have been evaluated, we discard solutions $x \in \mathcal{X}_j^*$ if $\hat{f}_i(x) - \hat{f}_i(x_{i-1}^*) > \lambda / j^{0.15}$. This threshold is used to remove candidates from the pool which no longer look promising. Note that $\lambda / j^{0.15}$ gets increasingly stringent so as to reduce the pool eventually to the optimal solution. To ensure convergence, a minimum amount of re-sampling is required throughout the procedure. We require for the number of evaluations $N(x)$ of any candidate x that $N(x) \geq \lceil j^{0.5} \rceil$, and determine further objective function values for x if the condition is not fulfilled.

At iterations i where $i \neq \lfloor j^{1.5} \rfloor$, we re-sample an existing candidate $x \in \mathcal{X}_j^*$ according to a pre-defined *re-sampling strategy*. We use an epsilon greedy approach: we sample with probability $\sigma(j) := 0.99/j^{0.5}$ a solution $x \in \mathcal{X}_j^*$ randomly, and with probability $1 - \sigma(j)$ we choose $x := x_i^*$ (where x_i^* is the best currently known average cost solution). We then evaluate $f(x, S)$ for a randomly chosen scenario S and update average objective function value $\hat{f}_i(x)$ and evaluation count $N(x)$ accordingly. We stop the procedure if we have evaluated all candidates in the current pool over N^{max} scenarios. We choose this stopping condition because a) we do not observe large variations in $\hat{f}(x)$ after a certain amount of evaluations, and b) there are often many very good candidates with similar $\hat{f}(x)$ in the pool (for large j), so reducing the pool further until only one candidate remains is not beneficial.

One advantage of the proposed framework is that only two parameters (λ and N^{max}) need to be set according to the problem at hand (see §3.5.1). However, the effectiveness and efficiency of the procedure depend on how well and how fast we can evaluate network cost $f(x, S) = G(x, S) + c^T x$ for any budget x and scenario S, which we discuss below.

3.4.2 Cost Evaluation for Given Solution and Given Scenario

To evaluate $G(x, S)$ for any candidate capacity budget x under traffic and capacity scenario S, we need to solve the integer program that underpins $G(x, S)$. However, this problem is $\mathcal{NP}$-hard since it represents an instance of the Multidimensional Multiple Choice Knapsack Problem (MMKP), as shown by Künnen and Strauss (2022). In fact, in order to determine an exact solution, we would need to solve an MMKP (i.e., the routing problem) for each of the $\prod_a |C^a|^{|U|}$ possible combinations of airspace configurations. This is infeasible for realistically-sized instances, and takes already 50 minutes for a small instance with 200 flights across five airspaces in one hour (see Künnen and Strauss (2022)). A potential approach to approximate $G(x, S)$ would be to solve its linear relaxation. However, we do not pursue this option since Starita et al. (2020) demonstrate on a large case study that it delivers poor results for this problem. Instead, we follow the approach in Künnen and Strauss (2022) and separate the sector-opening problem from the routing problem to estimate $G(x, S) \approx D(x, F^S, W^S)$ in polynomial time. More specifically, we conduct two steps to determine $D(x, F^S, W^S)$:

1. Determine best candidate configuration $C'(x, F^S, W^S)$ given the capacity budget x, traffic scenario F^S and uncertainty around capacity provision W^S;

2. Determine the routing of flights with lowest cost $D(C', F^S, W^S)$ given the candidate configuration C'.

Finding the best candidate configuration

In a first step, we want to determine candidate configuration $C' = \{c'_{au} : a \in A, u \in U\}$, which consists of individual configurations c'_{au} for each airspace and operating time. For this purpose, we first assign each flight to its shortest route, i.e., the route with zero displacement cost; a positive cost is incurred if a flight is "displaced" in time (delayed) or in space (re-routed). We assume that this traffic assignment (represented by allocation y) gives us a good indication of where capacity is required in the network. Given traffic assignment y, we can compute the capacity shortage k_{acu} for each airspace a, configuration $c \in C^a$ and time unit u. We define capacity shortage as the number of flights that exceed sector capacity limits, i.e., $k_{acu} := \sum_{l \in L^c} \left(\sum_{e \in E^l} \sum_{f \in F^S} \sum_{r \in R^f} b_{freu} y_r^f - \kappa_l^S \right)^+$, where $x^+ := \max\{x, 0\}$. We then determine configuration c'_{au} for each airspace a and time u as the feasible configuration with the lowest total capacity shortage by solving the following 'configuration integer linear program' (CILP):

$$\textbf{(CILP)} \quad \min_z \sum_{a,c,u} k_{acu} z_{acu} \tag{3.7}$$

$$\text{s.t.} \sum_{u \in U} \sum_{c \in C^a} \bar{h}_{ac} z_{acu} \leq x_a^S \qquad a \in A \tag{3.8}$$

$$\sum_{c \in C^a} z_{acu} = 1 \qquad a \in A, u \in U \tag{3.9}$$

$$z_{acu} \in \{0, 1\} \qquad a \in A, c \in C^a, u \in U. \tag{3.10}$$

In particular, c'_{au} is given by the configuration for which $y_{ac'u} = 1$ for each airspace a and operating time u. To ensure that the problem is always feasible with regards to constraint (3.8), we require that $x_a^S \geq |U|$ for all scenarios S and airspaces a because at least one sector needs to be operated at each time period. In the numerical experiments, we implement an even stricter lower bound for the airspace capacities, see §3.5.1.

It is easy to see that the CILP decomposes by airspace; we denote this decomposition by CILP-d. Even after decomposition, the resulting problem still represents a multiple choice knapsack problem (MCKP), which is $\mathcal{NP}$-hard. However, as there are typically no more than tens of configuration options per ACC in Europe in each time unit, the problem can be solved exactly in reasonable time. If larger airspaces need to be accounted for, or if the number of operating time units make the problem intractable, we could revert to heuristic approaches for the MCKP, such as the one proposed in Pisinger (1995).

To model capacity sharing, the approach outlined above for determining C' needs to

be adjusted to allow for sharing of capacities among two or more airspaces. Recall that $a \in A_g \subseteq A$ ($g \in G$) represent the airspaces in alliance g among which capacities can be shared. If capacity sharing is not permitted in certain airspaces, we combine these airspaces in alliance g' and set $x^0_{g'} := 0$. We can obtain the sector opening scheme under capacity sharing by solving the following mixed integer linear program (which we could decompose by sharing alliance g):

$$\textbf{(XCILP)} \quad \min_{h^0,z} \sum_{a,c,u} k_{acu} z_{acu} \tag{3.11}$$

$$\text{s.t. } (3.9), (3.10)$$

$$\sum_{u\in U} \sum_{c\in C^a} \bar{h}_{ac} z_{acu} \leq x^S_a + h^0_a \qquad a \in A \tag{3.12}$$

$$\sum_{a\in A_g} h^0_a \leq \sum_{a\in A_g} x^0_a \qquad g \in G \tag{3.13}$$

$$h^0_a \in \mathbb{N}^+ \qquad a \in A. \tag{3.14}$$

Determine lowest displacement costs with given configuration

After having determined configuration set C', we apply the MMKP-based heuristic proposed in Künnen and Strauss (2022) to determine the routing of flights with lowest displacement costs $D(C', F^S, W^S)$. The approach is summarized in Algorithm 5 in the Appendix. We initialize the routing procedure by assigning each flight $f \in F^S$ to the route with lowest displacement costs. Let $L' = \{l \in L^{c'} : c' \in C'\}$ be the sectors defined by C'. To establish a feasible solution, we then iteratively reassign flights on the most congested sector l^* until all sectors $l \in L'$ are within capacity limits κ^S_l (which depend on W^S). To decide which flight to reassign to another route, we compute a decision parameter γ^f_r that weighs the change in displacement costs with the change in capacity usage on the most congested sector. Finally, we test if we can use potential spare capacities to further improve this feasible solution. For that purpose, any flight f and route r is reassigned from current route r' to r, if this reassignment improves displacement costs while keeping the routing feasible.

As shown in Moser et al. (1997), Algorithm 5 has complexity $\mathcal{O}(m(n-o)^2 + mn)$, where $m = |L'|$ is total number of sectors given configuration C', $n = \sum_{f\in F^S} |R^f|$ is total number of flight-route combinations, and $o = |F^S|$ is total number of flights. Given this complexity, evaluating $D(C', F^S, W^S)$ is still computationally very expensive for larger networks and traffic scenarios. Therefore, we use Algorithm 5 to generate a large number of observations $\hat{D}$ given certain capacity budgets, traffic scenarios and capacity uncertainties, and apply a

linear regression to approximate D. In particular, we estimate D based on the capacity shortages $k_{ac'u}$ that a certain capacity budget and traffic and capacity scenario generates:

$$\hat{D}(C', F^S, W^S) = \beta_0 + \sum_{a \in A} \beta_a \sum_u k_{ac'u}. \tag{3.15}$$

The idea is that the number of capacity shortages directly influences the required amount of re-allocations of flights to alternative routes, and thus the size of total displacement costs. This choice of explanatory variable also offers another benefit: Since the CILP is decomposable by airspace, the solution with minimum capacity shortage k_{acu} will also be the one with minimum cost estimate (based on (3.15)). However, the XCILP can only be decomposed by alliance g (not by airspace a), because we allow capacities to be shared among airspaces in an alliance. Therefore, we may encounter cases where the solution with minimum capacity shortage does not provide minimum costs. To avoid this, we use parameters $(\beta_a)_{a \in A}$ as weights in the objective function of the XCILP to directly optimize over costs. In particular, we replace function (3.11) in the XCILP to get:

$$\begin{aligned} \textbf{(XCILP*)} \quad & \min_{h^0, z} \sum_{a,c,u} \beta_a k_{acu} z_{acu} \\ & \text{s.t. } (3.9), (3.10), (3.12), (3.13), (3.14). \end{aligned}$$

We denote the decomposition by alliances of the XCILP* by XCILP-d*. Overall, the CILP-d (or XCILP-d* for capacity sharing) together with regression (3.15) provide us with cost estimate $\hat{D}$ for any given capacity budget and scenario. We use this estimate in each iteration of Algorithm 2 to evaluate $f(x, S) \approx \hat{D}(C', F^S, W^S) + c^T x$. Since we can separately determine estimates $\hat{D}$ for each airspace (or alliance in the case of capacity sharing), we also run Algorithm 2 individually for each airspace (or alliance), which significantly speeds up the solving process. To further improve estimates $\hat{D}$, we could re-run Algorithm 5 and update parameters β after a fixed amount of iterations, using only budgets x that are close to the current optimum. However, we decide not to apply this updating procedure because (a) re-applying Algorithm 5 consumes sizable computational resources, and (b) the variation in displacement costs across scenarios is too large to warrant a more exact cost estimate per scenario.

3.5 Numerical Results

The objective of our research is two-fold: (1) We want to assess the performance of our stochastic optimization approach in finding optimal strategic capacity levels, and (2) to guide decision makers towards an operable and effective design of capacity sharing. In order to assess the value of capacity sharing, we need a case study of sufficient size (in terms of time horizon and number of airspaces) such that capacity sharing services may actually be demanded. For this purpose, we use a case study based on real flight data covering a 6-hour time period on the busiest day of 2016, see §3.5.1. We discuss results of our experiments on our stochastic approach and on capacity sharing in §3.5.2 and §3.5.3, respectively.

3.5.1 Case study description

We use the network data based on the large case study in Starita et al. (2020) and increase the considered time horizon and the number of flights. The real data set covers large parts of en-route airspace in Western Europe and was obtained using Eurocontrol's Demand Data Repository (DDR2) service. On the capacity side, the case study includes 15 ACCs (i.e., airspaces) across 8 ANSPs, which consist of 177 elementary sectors in total. These elementary sectors are combined in various ways to form a total of 173 different configurations, which were selected among the most frequently used ones in 2016. Figure 3.5.1 shows exemplary scheduled traffic across the considered network at a specific point in time. The capacity costs used in the case study are average costs per sector-hour for each ANSP, computed based on Eurocontrol (2018b). We treat these costs as variable costs in the simulation because expanding capacity will make it necessary to hire additional ATCOs in the long run (and vice versa). To define the solution space $\mathcal{X}$, we determine the minimum sector-hours $\underline{x}_a$ and maximum sector-hours $\overline{x}_a$ ($a \in A$) that each airspace operated based on historic data; we then have $\mathcal{X} = \left\{x \in [\underline{x}_a, \overline{x}_a]_{a \in A}, x \in \mathbb{N}^{|A|}\right\}$.

On the demand side, the data includes all flights in the selected network on September 9, 2016, the busiest day in European airspace that year. For the simulation, we restrict the time frame to 9am – 3pm and select the 2,400 scheduled flights crossing the network during this time. Among the remaining flights in the data set, we randomly choose 2,500 to serve as a pool of non-scheduled flights. To create the traffic scenarios, we uniformly sample from this pool the non-scheduled flights, where the number of flights is drawn from a normal distribution with mean of 600 (to generate on average 20% non-scheduled traffic) and standard deviation of 200. Each traffic scenario F^S then consists of all 2,400 scheduled

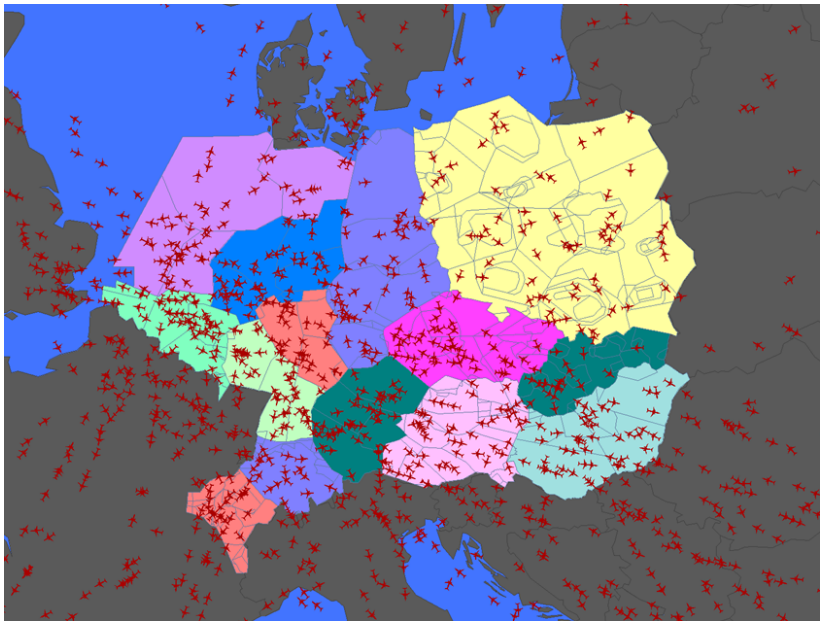

Figure 3.5.1: Snapshot of analysed flights at a specific point in time.
Each colour represents one airspace (lower airspace in Czech Rep. cannot be shown in 2D graph).

flights and a sample of non-scheduled traffic, giving on average 3,000 total flights.

Each flight in the data set has a reference (shortest) route with zero displacement cost and up to 12 different alternative routing options (in the horizontal or vertical plane). Reference and alternative routes were generated using Eurocontrol's Network Strategic Tool (NEST) based on the last filed flight plan for every flight in a data set. In particular, we identify up to 40 horizontal route options for each flight, and then compute the vertical profile (considering both the actual route network and constraints in the horizontal and vertical planes). Among the resulting route options, the shortest trajectory is then chosen based on lowest total distance; wind conditions are not considered since the capacity decision is made in the strategic phase of ATM in which weather forecasts are still unreliable. We assume that a flight can only be subject to one demand management measure: either delay or re-routing (this also keeps the number of total route alternatives low). Therefore, we add three potential delay options to the shortest route of each flight: 10, 20 and 30 minutes. To further reduce solution time of the model (especially Algorithm 5), we pre-process the routes of all scheduled flights and keep only frequently used route options. Displacement costs are computed for each flight and route option separately and consist of re-routing and delay costs. Re-routing costs include mainly additional fuel costs, crew and passenger costs and are estimated based on Cook and Tanner (2015) and Eurocontrol (2018a). Delay costs depend to a large extent

on the AU's business model. Thus, let ω represent a scaling factor for delay cost of each flight based on the AU's business model $v(f)$: 0.4 for charter flights and low cost carriers, 1 for full-service carriers, 1.75 for business aviation and 10 for flights which are typically exempt from demand management measures. Also, let $\bar{d}_r^f$ represent delay cost references based on Cook and Tanner (2015) which differ by aircraft type and increase non-linearly with the duration of delay. We can then estimate delay costs d_r^f for all flights f and delay options $r \in R_{delay}^f$ as follows:

$$d_r^f := \omega_{v(f)} \bar{d}_r^f \qquad r \in R_{delay}^f, f \in F.$$

To model a real-life setting, we need to account for both demand- and capacity-side uncertainties. Demand-side uncertainty stems from non-scheduled flights (in terms of total number of such flights and where and when they appear in the network), which is modeled through our traffic scenarios F^S described above. For capacity-side uncertainty W^S we distinguish between internal (i.e., ATCO shortages) and external (i.e., adverse weather) effects. This distinction is particularly important because internal effects can be mitigated, at least partly, through capacity sharing while external effects cannot. To estimate internal effects, we analyze air traffic flow management (ATFM) regulations due to ATCO staffing reasons in the considered network. An ATFM regulation is issued when demand exceeds capacity in an airspace volume for a period of time. We partially rely on analysis of 2016 ATFM regulation data to derive frequencies of staffing regulation occurrences for each airspace in the case study, which we show in the Appendix. External effects can sometimes severely reduce capacities in an affected airspace for a certain amount of time, which is why we apply a more differentiated probability distribution here. Again, based on historical ATFM regulation data, we assume that the capacity of any one elementary sector (and the collapsed sector containing it), is reduced to 90% in 10% of cases, to 70% in another 10% of cases and to 50% in 5% of cases.

Two parameters in Algorithm 2 (λ and N^{max}) need to be tailored to our case study. Parameter λ determines which candidates we accept into the pool and which ones we discard (after a certain amount of evaluations). We decide to only accept candidates that stay within 20% of the current minimal cost solution. Accordingly, we set $\lambda :=$ 2,000 when running Algorithm 2 for each airspace, and $\lambda :=$ 4,000 when running it for each alliance (if capacity sharing is allowed). Furthermore, we set $N^{max} :=$ 300 since we do not observe significant changes in $f(x, S)$ if we evaluate more than 300 scenarios.

3.5.2 Value of the stochastic approach

Simulation setting and evaluation

To analyze whether the stochastic approach leads to better capacity planning decisions, we compare it with a deterministic benchmark, which we describe below. The quality of a capacity management decision is assessed based on the expected network costs (both displacement and capacity costs) that it creates. We test the approaches on the setting without capacity sharing, i.e., all airspaces can only use their own, local capacities to manage traffic demand. In the stochastic approach we determine optimal capacity budget $x^* = (x^*_a)_{a\in A}$ using Algorithm 2 as discussed in §3.4. Here, we use a fixed set of 300 training scenarios (stored in $\mathcal{S}_1$) for traffic and capacity uncertainty on which each candidate budget is evaluated.

To establish a benchmark for our stochastic approach, we use the strategic capacity planning model proposed by Starita et al. (2020). In this approach, the optimal budgets x^S are determined separately for each scenario $S \in \mathcal{S}_1$ by decomposing the problem underpinning $G(x, S)$ into a master- and sub-problem. After evaluating all scenarios, the benchmark capacity budget x^B is determined using the so-called risk-based policy, which performed best in Starita et al. (2020). More specifically, we have $= x^B = (x^{S'}_a)_{a\in A}$ where each $x^{S'}_a$ is chosen among all x^S_a ($S \in \mathcal{S}_1$) such that the probability that any other budget x^S_a (for $S \neq S'$) has higher capacity in airspace a is less than a given ϵ (which is set to 0.05).

Once the capacity budgets x^* and x^B have been determined, we assess their performance on a set of 200 testing scenarios (stored in $\mathcal{S}_2$, and sampled from the same distributions as $\mathcal{S}_1$). When applying Algorithm 5 to solve the routing problem, we implicitly assume that the NM can autonomously decide on the routing of each flight (in the strategic phase). However, on the day of operation the decision is made iteratively between the NM and airspace users (AUs). Therefore, we use a discrete-time event simulation to test each capacity decision, in which the NM decides on the sector-opening scheme of each airspace (using the CILP), communicates all feasible trajectories to the AU and the AU then decides on the final trajectory among all feasible options. To model the AU decision, we assume that the AU chooses the trajectory with lowest AU-adjusted displacement cost δ, which we determine as follows:

$$\delta^f_r := \begin{cases} d^f_r & \text{if } r \text{ is a re-routing option,} \\ N(d^f_r, 0.1) & \text{if } r \text{ is a delay option.} \end{cases}$$

Here, $N()$ denotes the draw from the normal distribution. We include this uncertainty for

delayed routes because, in contrast to rerouting costs, delay costs depend more strongly on AU-specific parameters. Note that we do not consider ATM service charges in the trajectory decision because we assume that charges are route-independent (for details on the so-called airport-pair charging principle, see Pavlović and Fichert (2019)).

Furthermore, we introduce capacity uncertainty dynamically in the discrete-time event simulation by updating the available capacity budget x^{*S} and sector capacities κ^S after every time period u. In contrast, the materialized demand is fixed for each run of the simulation (and thus not updated dynamically), because we assume that flights are fully known at the beginning of the day of operation. The procedure is summarized in Algorithm 3. In addition to existing notation, we denote by F_u the flights that depart within u, by S_u a scenario that models capacity uncertainties within u and by $f(r)$ the flight corresponding to route r.

Algorithm 3 Discrete-time event simulation to test capacity decision.

Input: Flights F, capacity decision x^*, scenarios $\mathcal{S}_2$
Initialize: Iteration counter $i := 0$, Routing $R_0^* := \emptyset$
for $u \in U$ **do**
  Sample new scenario $S_u \in \mathcal{S}_2$ and update capacity uncertainty: x^{*S_u} (ATCO shortages) and κ^{S_u} (adverse weather)
  Run CILP to determine best configuration $C'(x^{*S_u}, F, R_i^*, \kappa^{S_u})$
  for $f \in F_u$ (sorted chronologically by departure time) **do**
    Determine set of feasible routes:
$\bar{R}^f := \left\{ r \in R^f : \sum_{e \in E_l} \left(b_{f,r,e,u} + \sum_{j \in R_i^*} b_{f(j),j,e,u} \right) \leq \kappa_l^{S_u} \quad \forall a \in A, u \in U, l \in L^{c'_{au}} \right\}$
    Determine route choice based on adjusted displacement cost δ: $r^f := \arg\min_{r \in \bar{R}^f} \delta_r^f$
    Update sets: $R_{i+1}^* := R_i^* \cup r^f$, $i = i + 1$
  end for
end for
Output: Routing R_i^* with displacement cost $D(x^*, F^S, W^S) := \sum_{r \in R_i^*} d_r^{f(r)}$

Discussion of results

Before we discuss the performance of the stochastic and benchmark approaches to minimize network costs, we comment on efficiency and effectiveness of regression (3.15) to approximate displacement cost $G(x, S)$ for a given budget x and scenario S. For this purpose, we compare the cost estimates from regression (3.15) on 2,280 instances (each instance is a combination of budget and scenario) against costs determined through Algorithm 5 and the CILP. With a mean absolute error of 6,191 (or 6.4% of average displacement costs) and R^2 of 0.96, we confirm that regression (3.15) adequately approximates displacement costs for our purpose.

To illustrate the relationship, we show in Figure 3.5.2 the correlation between capacity shortages (aggregated across airspaces and time for purpose of illustration) and displacement costs based on Algorithm 5 for all instances. Applying regression (3.15) instead of Algorithm 5 (together with the CILP-d) to estimate displacement costs significantly reduces our solution time. For the case study at hand, we manage to reduce average run time for one instance from 20 minutes to under one second, including running the CILP-d.

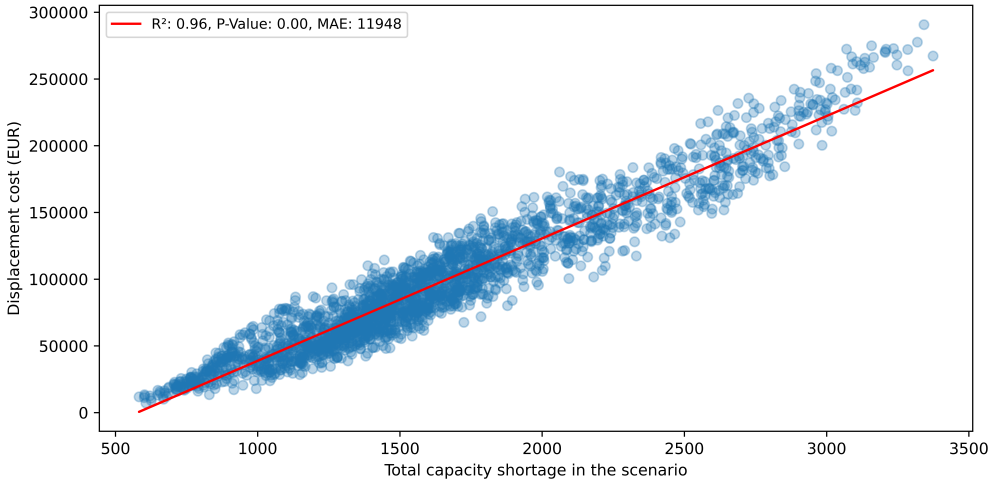

Figure 3.5.2: Correlation between total capacity shortage and disp. costs (n = 2280 instances).

Table 3.5.1 summarizes the simulation results for the capacity decisions x^* and x^{B*} of the stochastic and deterministic approach on 200 scenarios, without capacity sharing. The stochastically determined capacities lead to significantly lower network costs than the benchmark capacity decision, since the displacement cost savings of €118,213 (or 53%) outweigh the higher capacity costs. Overall, network costs are reduced by €96,834, or 28%. Figure 3.5.3 illustrates the performance of the stochastic vs deterministic capacity decision across 50 of the 200 evaluated scenarios. Not only does the stochastic solution create lower network costs in all 50 scenarios, but it also reduces the *variation* in cost (which is also reflected in the standard deviations in Table 3.5.1). This is important because it implies a more stable network performance and thus higher reliability of service provision to airspace users (e.g., airlines).

The results also show that the number of flights that cannot be assigned to non-dummy routes is 1.4% (on average across scenarios) for the stochastic approach, which shows that almost all flights can be routed within the given time frame and network conditions. For the benchmark, this share increases to 3.8%. Furthermore, under stochastically-determined

Table 3.5.1. Simulation results for solution approaches (n = 200 scenarios).

Approach	Cap. cost	Displ. cost	Network cost	Not assigned	Run time
Stochastic	140,936	106,112	247,048 ± 39,295	1.4%	126 min.
Benchmark	119,556	224,325	343,882 ± 67,460	3.8%	3017 min.
Difference	21,379	-118,213	-96,834 ± 36,558*	-2.4%	-2,891 min.
Difference (%)	+18%	-53%	-28%	-63%	-96%

* Significant at 95% confidence level.

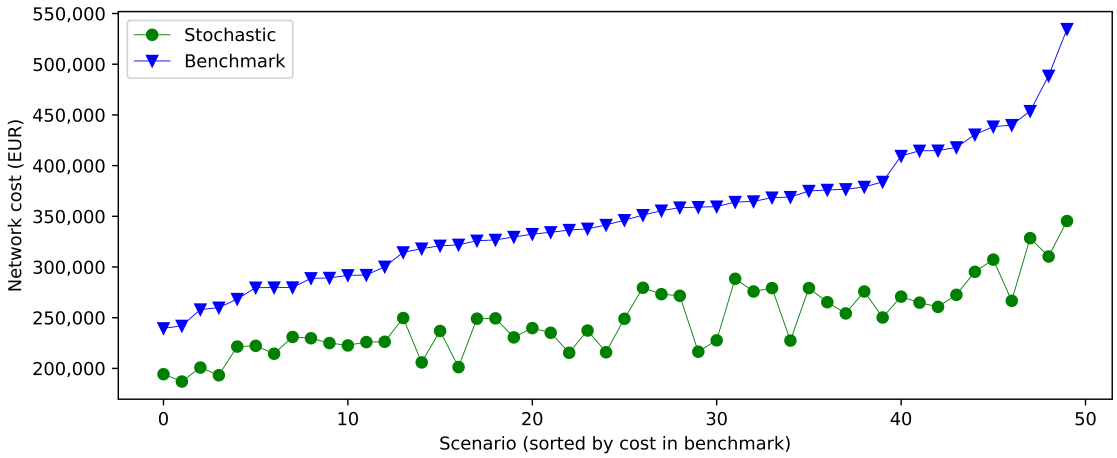

Figure 3.5.3: Network cost across 50 scenarios for stochastic and benchmark approach.

capacities a lower share of flights need to be rerouted or delayed than under benchmark capacities (see Figure 3.5.4). In fact, more than 84% of flights can take the shortest route for the stochastic approach vs 72% for the benchmark. These figures confirm that stochastic capacity planning can substantially reduce network cost and improve network performance.

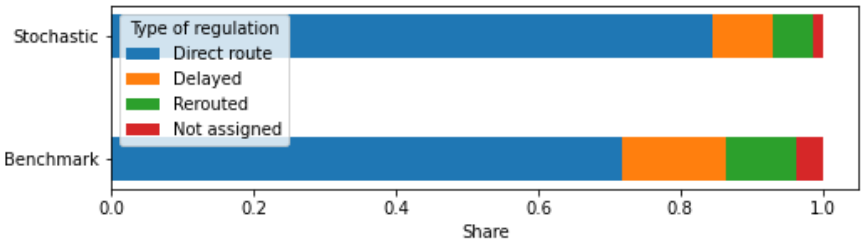

Figure 3.5.4: Regulations for stochastic vs benchmark approach (n = 200 scenarios).

Next to cost, an important consideration for the network manager (or ANSPs in the case of local decision-making) is the speed with which capacity decisions can be made. With the

proposed stochastic approach, we determine the capacity decision x^* in around two hours, and thus sufficiently fast for practical application. Figure 3.5.5 visualizes an exemplary path of Algorithm 2 to determine the capacity decision x^* (shown for airspace EDUUUTAS). We see that the procedure converges quickly towards a good solution, with only small variations in the recommended capacity level after around 700 iterations. Since we can decompose the problem by airspace, we could also parallelize the process by solving x_a^* for each airspace simultaneously, which would further reduce the run time to around 10 minutes. In contrast, the run time for the benchmark approach exceeds 3,000 minutes (or around 2 days), and the problem cannot be decomposed and parallelized which renders the method less practical in a dynamic environment. It is important to note that the solution time for the stochastic approach does not include the time to estimate parameter β for regression (3.15), which takes around 20 minutes for each of the roughly 2,000 observations of $G(x, S)$. However, this process can be parallelized for every observation, and it would only have to be conducted once for all future capacity decisions.

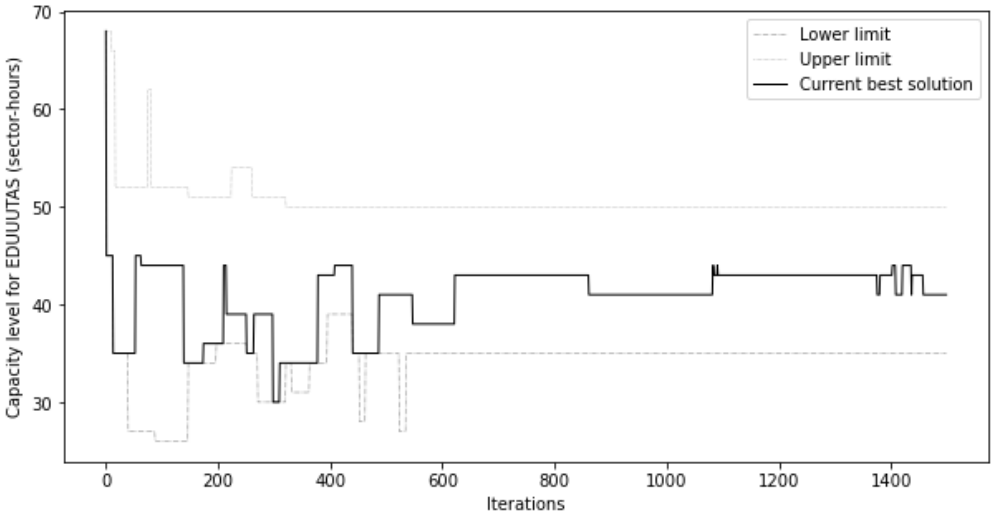

Figure 3.5.5: Exemplary path to determine capacity decision in stochastic approach.

Sensitivity Analysis

In order to test how robust the presented results are to changes in traffic intensity, we conduct a sensitivity analysis across three traffic levels: next to medium traffic with on average 3,000 flights used so far, we analyze performance under low traffic with 2,500 flights and high traffic with 3,500 flights. In each case, we continue to sample traffic scenarios F^S from the

data set with 80% scheduled and 20% non-scheduled flights. The results in Table 3.5.2 show that the stochastic approach outperforms the benchmark by €94,334 (or 36%) under low traffic and by €67,184 (or 15%) under high traffic. The improvement is lowest under high traffic because airspaces may simply not be able to increase capacities further to manage all traffic. Furthermore, between 3.4–4.5% of flights cannot be assigned to non-dummy routes under benchmark capacities; the values for the stochastic approach are much lower. Even though the run time for the stochastic approach increases to 156 minutes under high traffic, it stays within practical limits and well below the solution time for the benchmark. Finally, next to the benefits in solution quality and time, the proposed stochastic approach features a methodological advantage over existing deterministic approaches: Rather than fixing scenarios and finding appropriate capacities for each, it fixes the capacities first and then determines its performance across a multitude of scenarios. It is this feature that allows us to use the stochastic approach to determine 'optimal' capacity levels in the setting with capacity sharing among airspaces, which we discuss in the following.

Table 3.5.2. Simulation results for solution approaches across traffic levels (n= 200 scenarios).

	Network cost		**Not assigned**		**Run time (min)**	
Model/Traffic flow	Low	High	Low	High	Low	High
Stochastic	166,369	386,800	0.3%	3.3%	111	156
Benchmark	260,703	453,984	3.4%	4.5%	3,523	2,937
Difference	-94,334 ± 58,048	-67,184 ± 26,841*	-3.1%	-1.2%	-3,412	-2,781
Difference (%)	-36%	-15%	-92%	-27%	-97%	-95%

* Significant at 95% confidence level.

3.5.3 Value of capacity sharing

Capacity sharing concepts and evaluation

To provide guidance for decision makers on how to design a potential capacity sharing scheme in European ATM, we test three different design options:

1. Capacity sharing across airspaces within the same ANSP (*Cross-ACC*): Every airspace (i.e., area control center, or ACC) within an ANSP can leverage the shared capacity of its ANSP, next to its own capacity, to manage traffic demand.

2. Capacity sharing across ANSPs (*Cross-border*): In this setting, capacity can be shared between any combination of ANSPs, without restrictions due to location or technology.

3. Capacity sharing among ANSPs with the same technology provider (*Common tech*): Next to geographic proximity, the use of common technological infrastructure is a major criterion for sharing of capacities across ANSPs. Therefore, in this setting an alliance can only be formed among airspaces that use the same air traffic technology provider (the provider for each airspace is reported in Table 3.C.1 in the Appendix).

To model capacity sharing, we apply the (XCILP-d*) in each evaluation of Algorithm 2 to determine the best candidate configuration for each budget x and alliance g given solution space $\mathcal{X}_g$, where $\mathcal{X}_g = \left\{x \in [\underline{x}_a, \overline{x}_a]_{a \in A_g}, x^0 \in [\underline{x}^0_a, \overline{x}^0_a], x, x^0 \in \mathbb{N}^{|A_g|}\right\}$. Parameters $\underline{x}_a, \overline{x}_a$ for $a \in A$ are given by the case study, see §3.5.1. In order to fulfill condition (3.1) for budgets x, we set $\underline{x}^0_a(x) = (\bar{x}_a - x_a)^+$ and $\overline{x}^0_a(x) = \overline{x}_a - x_a$ for $g \in G$. The definition of G and $A_g (g \in G)$ varies for each of the three design options above. For the *cross-ACC* setting, we define each ANSP as a separate alliance among which capacities can be shared. For the *cross-border* setting, the definition of alliances is less straight-forward. Since airspaces can be combined freely in this setting to form an alliance, the number of potential set partitions increases exponentially with the number of airspaces. Therefore, we restrict the size of a cross-border alliance to up to two airspaces, and use a structured approach to determine which combination of such alliances (i.e., set partitions) promises to deliver the best cost performance. Let $\mathcal{G}$ contain all 1- and 2-tuples of airspaces $a \in A$, i.e. $\mathcal{G} := \{a \in A\} \cup \{(a,b) | a \in A, b \in A, a \neq b\}$, and let $\mathcal{G}_a$ ($a \in A$) be the subset of $\mathcal{G}$ that includes airspace a, i.e. $\mathcal{G}_a := \{g \in \mathcal{G} | a \in A_g\}$. To judge the cost performance of an alliance, we use Algorithm 2 to determine capacity decision $x^*_g = (x^*_a)_{a \in A_g}$ with network costs $\hat{f}(x^*_g)$ for each $g \in \mathcal{G}$. We can then determine the set partition G for the *cross-border* setting with the following integer program:

$$\min_{\theta} \sum_{g \in \mathcal{G}} \theta_g \hat{f}(x^*_g) \tag{3.16}$$

$$\text{s.t.} \sum_{g \in \mathcal{G}_a} \theta_g = 1 \qquad a \in A \tag{3.17}$$

$$\theta_g \in \{0,1\} \qquad g \in \mathcal{G} \tag{3.18}$$

Decision variable θ_g determines whether alliance g is used in the final set partition G; we have $G = \{g \in \mathcal{G} | \theta_g = 1\}$. The objective function (3.16) minimizes the total expected network costs, and constraint (3.17) ensures that each airspace is contained in the final set

partition. To determine the alliances for the *common tech* setting we again solve the integer program (3.16)–(3.18), but we restrict $\mathcal{G}$ to those alliances in which both airspaces use the same provider. The set partitions used for all three settings are summarized in Table 3.C.3 in the Appendix.

Apart from the alliances, a further consideration in designing capacity sharing is the marginal cost for such sharing services. We require that the cost per sector-hour of providing capacity sharing is strictly higher than the local cost for ATM services, because ATCOs will need to be reimbursed for the additional qualification and training associated with this task. We set $c_a^0 := (1+\rho)c_a$ for all $a \in A$ and require for the cost markup $\rho > 0$. In particular, we set $\rho := 0.1$ in the baseline and compare different markups in a sensitivity analysis.

As discussed in §3.5.2, we assess the performance of each of the proposed capacity sharing settings with a two-staged simulation study: We first determine budget x^* using Algorithm 2 (using 300 training scenarios in $\mathcal{S}_1$), and test its performance using the XCILP* and Algorithm 3 (on a new set of 200 testing scenarios in $\mathcal{S}_2$). Again, simulations are run using AWS Batch with 4 CPU and 30GB RAM.

Discussion of results

Table 3.5.3 compares the cost performance in the simulation of the standard setting (without capacity sharing) with the three capacity sharing options, based on the network described in §3.5.1. We find that capacity sharing may reduce total network costs by €4,900 to €20,700 (around 2.0% to 8.4%) against the standard setting. Most notably, the total saving is realized with only small increases in capacity costs: In the *cross-border* setting, an investment of €430 (or 0.3%) in capacity costs leads to €8,248 (or 20%) lower displacement costs. The results also show that capacity sharing within the same ANSP (*cross-ACC*) delivers in fact larger savings than sharing among different ANSPs with the same technology provider (*common tech*). Since *cross-ACC* sharing would also be easier to implement, both politically and operationally, it should be the preferred setup. Note that one reason why we observe larger savings for *cross-ACC* sharing is that we allow alliances to contain up to four airspaces (since the German ANSP splits into four regions), while in *common tech* we allow only two.

The sector-hours reported in Table 3.5.3 show that out of the 912 total hours used in the standard setting, only a small fraction of 24 to 37 hours (or 2.6% to 4.1%) would need to be deployed virtually in case of capacity sharing. This shows that the benefits from capacity sharing may be reaped with a relatively small amount of cross-border resources.

As expected, the solution times with capacity sharing are somewhat longer than for the

Table 3.5.3. Simulation results of capacity sharing (n = 200 scenarios).

Setting	Cap. cost	Displ. cost	Netw. cost	Savings	Sector-hours	Run time
No sharing	140,936	106,112	247,048	-	912	126 min.
Cross-ACC	141,367	97,864	239,231	7,817 (3.2 %)	887 + 25	127 min.
Cross-border	141,413	84,911	226,324	20,723 (8.4 %)	875 + 37	173 min.
Common tech.	141,255	100,884	242,139	4,908 (2.0 %)	888 + 24	163 min.

standard setting, because the search space $\mathcal{X}$ becomes much larger. However, the run time stays below three hours in all settings, which allows practical implementation. Also note that the reported run times represent non-parallelized processes. By decomposing the problem, we can parallelize the solution process and solve it simultaneously for each airspace (or alliance, for capacity sharing). This would reduce the run time to around 8 minutes in the standard setting and 18-25 minutes for capacity sharing.

The savings in all settings are generated exclusively by reducing the average displacement costs rather than by reducing capacities. This is particularly relevant because lower re-routing costs also imply lower greenhouse gas emissions from reduced fuel burn. If an additional cost for emissions was considered, this would further increase the potential benefit reaped through capacity sharing. Figure 3.5.6 shows how the displacement costs vary for 50 of the 200 scenarios across the settings. We find that capacity sharing can reduce displacement costs especially when the costs in the standard setting are high. While costs under *No sharing* exceed €125,000 in 15 of the 50 scenarios (corresponding to cases with major disruption), the costs with *cross-border* sharing stay below this value for all but one scenario. This shows that the benefit of capacity sharing is especially large in case of major disruptions (e.g., due to adverse weather), in which case the displacement cost saving can amount to €30,000 or more. We also compare the performance to a lower bound on displacement costs, which we obtain by setting all capacities to their maximum historic reference values, i.e., $x_a^* := \overline{x}_a$ for all $a \in A$. We find that *cross-border* sharing realizes large parts of the total displacement cost reduction potential (on average over 68%) between the standard setting and the lower bound.

We also want to analyze which flights in particular benefit from the flexibility gained through capacity sharing. For this purpose, we compare the average displacement costs per flight for scheduled vs non-scheduled flights and different aircraft sizes between the standard setting and *cross-border* sharing (across five selected scenarios), see Table 3.5.4. The matching of aircraft types (e.g., Boeing 737) to small, medium or large aircraft size is shown in

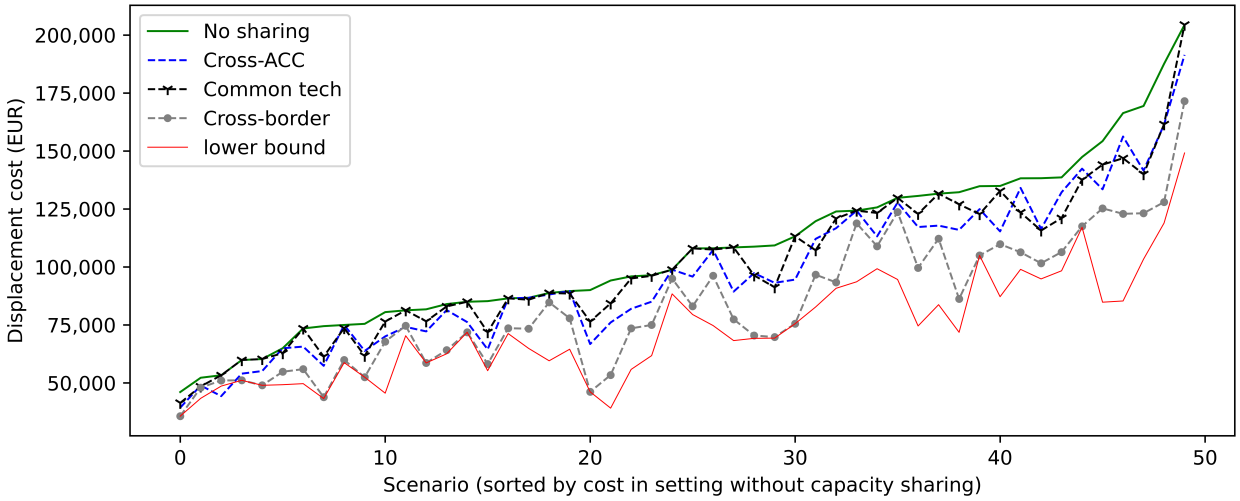

Figure 3.5.6: Displacement cost for capacity sharing settings in 50 scenarios (medium traffic).

Table 3.C.2 in the Appendix. We find that flights with small aircraft benefit more from capacity sharing (-43.7%) than flights with medium- to large-sized aircraft (-14.1% to -15.1%). This is likely because without capacity sharing, small aircraft are re-routed or delayed more heavily (in the case of capacity shortages) because their displacements are less costly than those of medium or large-sized aircraft; with capacity sharing, these displacements often become unnecessary, which disproportionately benefits small aircraft. Furthermore, we find that non-scheduled flights benefit more from capacity sharing (-28.4% in displacement costs) than scheduled flights (-14.1%). This is because non-scheduled flights are more likely than scheduled flights to operate on small aircraft, which benefit more from capacity sharing. Note also that the absolute displacement costs per scheduled flight is higher than for non-scheduled flights because scheduled flights are more likely to operate on large aircraft.

Table 3.5.4. Avg. displacement cost (EUR) per flight by setting (n = 5 scenarios).

	Aircraft type			**Flight type**	
Setting	Small	Medium	Large	Scheduled	Non-Scheduled
No sharing	47.7	31.2	85.7	38.9	27.9
Cross-border	26.9	26.5	73.6	33.4	20.0
Difference %	-43.7%	-15.1%	-14.1%	-14.1%	-28.4%

Finally, in Figure 3.5.7 we illustrate how capacity sharing helps increase the flexibility with which capacities can be adjusted within an alliance. The capacities used by the two Swiss airspaces (LSAZUTA and LSAGUTA) change frequently across the 50 scenarios, to

adjust to the materialized demand and capacity provision.

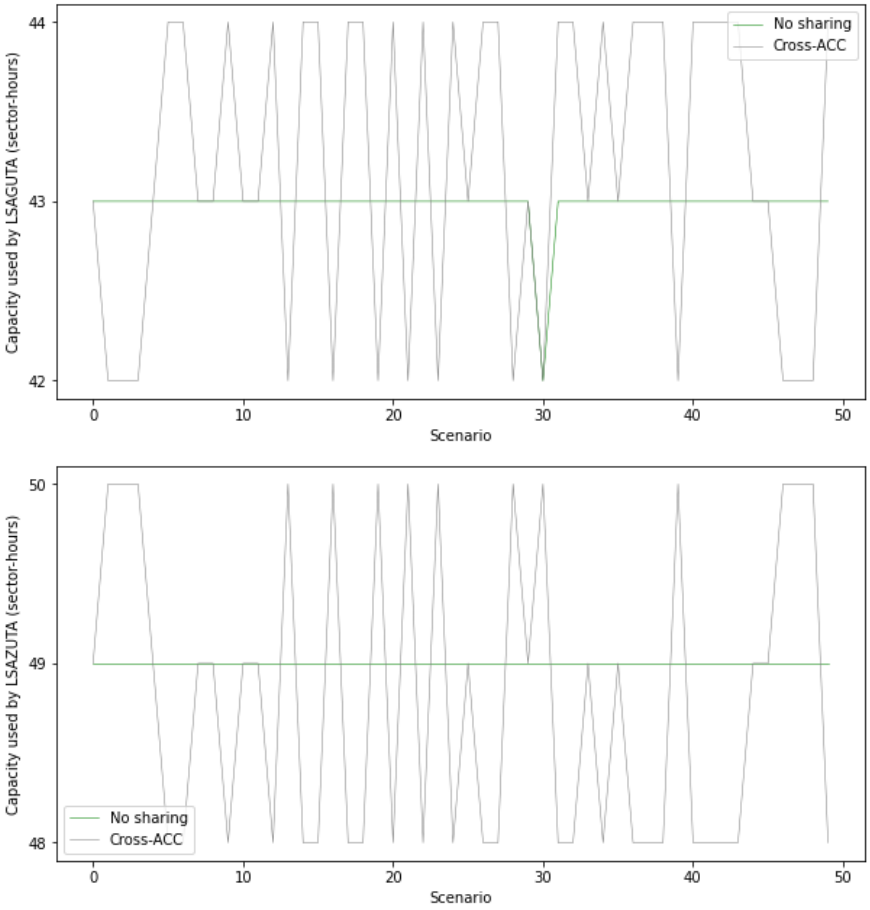

Figure 3.5.7: Capacity used by airspaces in Swiss ANSP across 50 scenarios (in *cross-ACC*).

Sensitivity Analysis

While capacity sharing can work to improve network performance, this comes at a cost. For that purpose, we want to determine what influence the marginal cost of capacity sharing has on the value of such a service. We conduct a sensitivity analysis with increased cost markups of $\rho := 0.2$ and $\rho := 0.4$ for the *cross-border* setting (which performed best in the tests above), see Table 3.5.5. Since the number of sector-hours required for capacity sharing is rather small (see Table 3.5.3), the total capacity costs increase only slightly if capacity sharing comes at a higher markup. In particular, total savings (compared to *No sharing*) decrease only from €20,724 to €20,247 (€19,293) for a markup of 20% (40%). This shows that if capacity sharing is implemented with only the few cross-border resources that we show

are needed, then the marginal cost of such sharing services will not play a dominant role.

Table 3.5.5. Sensitivity of *cross-border* capacity sharing to cost (n = 200 scenarios).

Setting	Cap. cost	Displ. cost	Network cost	Savings
No sharing	140,936	106,112	247,048	-
Cross-border (10%)	141,413	84,911	226,324	20,724 (8.4%)
Cross-border (20%)	141,890	84,911	226,801	20,247 (8.2%)
Cross-border (40%)	142,844	84,911	227,755	19,293 (7.8%)

As in §3.5.2, we also test how well capacity sharing performs under different traffic intensities, see Table 3.5.6. We find that *cross-border* sharing delivers the best results under low traffic, but *cross-ACC* sharing performs best under high traffic. Furthermore, the *common tech* setting shows comparatively low performance across all three traffic intensities. On our case study, this provides some evidence that establishing a shared technological infrastructure within alliances presents an important precondition to reap the benefits from capacity sharing. Note, however, that the savings from capacity sharing are not significant at 95% under either low or high traffic.

Table 3.5.6. Sensitivity of capacity sharing to traffic levels (n = 200 scenarios).

Traffic level	**Low**		**High**	
	Network cost	Savings	Network cost	Savings
No sharing	166,369	-	386,800	-
Cross-ACC	163,144	3,225 (1.9%)	**380,862**	5,938 (1.5%)
Cross-border	**151,845**	14,525 (8.7%)	381,198	5,601 (1.4%)
Common tech.	162,722	3,647 (2.2%)	383,401	3,398 (0.9%)

Finally, Figure 3.5.8 compares the displacement cost performance of the standard vs *cross-border* setting for all three traffic levels and across 50 of the 200 evaluated scenarios. It is easy to see that the maximum potential to reduce displacement costs is lowest under low traffic and highest under high traffic (as reflected in the gap between *No sharing* and the lower bound). However, while *cross-border* sharing reduces costs to the lower bound in most of the scenarios under low traffic, it fails to do so under high traffic. This is likely because if traffic volumes exceed a certain level, all ATCOs will be required to manage local traffic and thus will not be available for capacity sharing. The results provide guidance under which conditions capacity sharing adds value: the expected traffic level needs to be large enough

so that substantial displacements occur, but small enough so that airspaces do not require all of their ATCOs themselves.

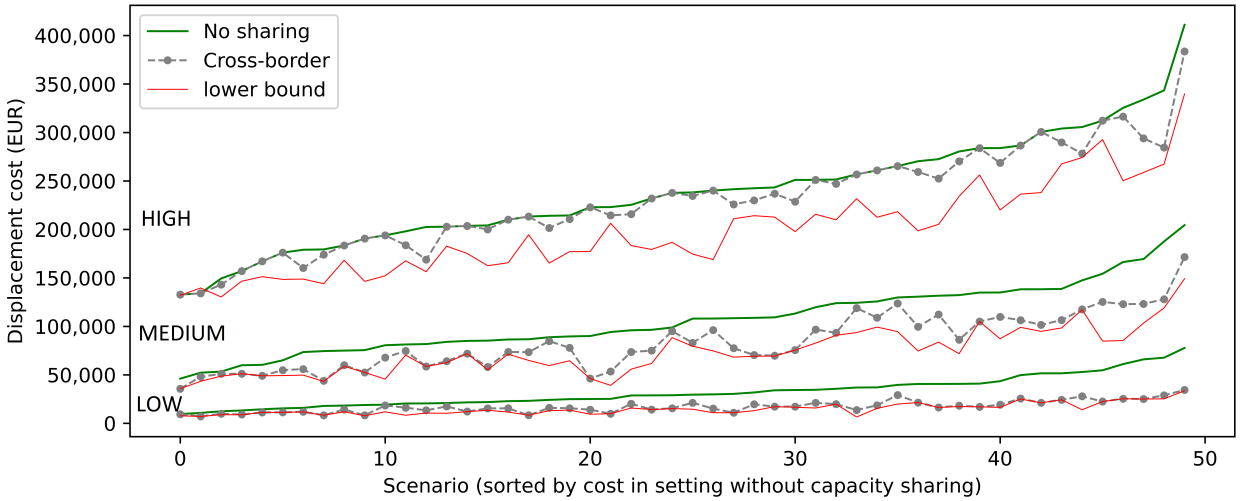

Figure 3.5.8: Disp. cost saving of capacity sharing across traffic levels (n = 50 scenarios).

3.5.4 Limitations

The proposed methodology features some limitations. Firstly, while Algorithm 2 is scalable in the number of airspaces in the network, it does not scale well in the number of airspaces contained in each alliance (if capacity sharing is allowed). This is because we cannot decompose XCILP* by airspace but only by alliance. Secondly, to determine the regression parameters required for Algorithm 2, we need to run Algorithm 5 to generate a sufficiently large number of displacement cost observations. This procedure works well for the case study at hand, but does not scale well for much larger networks. Thirdly, we base our regression on estimates of displacement costs obtained from a heuristic procedure, which can only approximate these costs; this limits the quality of cost estimates obtained through regression (3.15).

Furthermore, there are a few assumptions made in the evaluation that may limit the generalizability of results. On the one hand, we assume that more sector-hours will lead to higher total ATCO costs in the long-term. However, this effect may only materialize inflexibly, since ATCOs are hired and trained years before becoming operational. On the other hand, we use a time horizon of 6 hours on one day to make strategic decisions (with regards to capacity levels and capacity sharing concepts), whereas in practice these decisions will need to be judged based on the performance over a sustained period of time. Finally,

further practical considerations such as rostering practices for ATCOs are neglected in our simulation, but may have an impact on how well the strategic decisions translate into better network cost performance.

3.6 Conclusions

In European ATM, we observe large demand-capacity imbalances due to static capacities paired with large uncertainties in demand and capacity provision. To reduce the impact of these imbalances, we propose a capacity sharing scheme in which capacities can be shared among airspaces in an alliance. We develop an efficient approach to determine capacity budgets for each airspace that perform well across various scenarios (of demand and capacity uncertainty). A simulation study on a realistically-sized network with around 3,000 flights shows that the stochastically-determined capacity decision significantly outperforms a deterministic benchmark. As additional benefit, stochastically-determined capacities can lead to a more stable network performance by reducing variation in displacement costs. We use the proposed methodology to test different settings for capacity sharing, and find in the case study that cross-border sharing can reduce network costs by up to 8.4%. The reduction potential does however depend on the design of the alliances and the traffic intensity. If alliances can only be formed among ANSPs that use the same technological infrastructure, the cost savings reduce to 2–3% in our case study. Furthermore, if traffic volumes increase beyond a certain threshold, the benefits from capacity sharing decline because each airspace requires their local ATCOs themselves.

Author Statement

Jan-Rasmus Künnen: Writing - Original Draft, Software, Methodology, Formal analysis, Investigation. **Arne Strauss**: Supervision, Conceptualization, Methodology, Writing – Review & Editing, Funding acquisition. **Radosav Jovanović**: Conceptualization, Funding acquisition, Writing – Review & Editing. **Nikola Ivanov**: Conceptualization, Data Curation, Writing – Review & Editing. **Frank Fichert**: Conceptualization, Funding acquisition, Writing – Review & Editing. **Stefano Starita**: Software.

Acknowledgements

This project has received funding from the SESAR Joint Undertaking within the framework of SESAR 2020 and the EU's Horizon 2020 research and innovation programme under the Grant Agreement Number 893380. The opinions expressed herein reflect the author's view only. Under no circumstances shall the SESAR Joint Undertaking be responsible for any use that may be made of the information contained herein.

Appendix

3.A Latin Hypercube Sampling

This section describes the Latin Hypercube Sampling procedure used to build random capacity samples, see Algorithm 4. Each iteration of this algorithm generates $\bar{\eta}$ different samples, stored into a $\bar{\eta} \times |A|$ matrix, X. Let us also define $\mathcal{X}^a = \{\bar{x}_a^{min}, \bar{x}_a^{min} + \epsilon, \bar{x}_a^{min} + 2\epsilon, ..., \bar{x}_a^{max}\}$, with $\epsilon = (\bar{x}_a^{max} - \bar{x}_a^{min})/\bar{\eta}$. Values $\bar{x}_a^{min}$ and $\bar{x}_a^{max}$ represents the minimum and maximum number of sector hours for airspace a. These values are defined by the structure of the airspace's configurations and the length of the time horizon.

Algorithm 4 Latin Hypercube Sampling

for $\eta = 1$ to $\bar{\eta}$ **do**
  $\mathcal{A} \Leftarrow A$
  while $\mathcal{A} \neq \emptyset$ **do**
    a is uniformly drawn and removed from $\mathcal{A}$
    $\bar{x}_a$ is uniformly drawn and removed from $\mathcal{X}^a$
    $X[\eta][a] = \bar{x}_a$
  end while
end for
return X

3.B Heuristic for routing problem

Algorithm 5 MMKP-based heuristic for routing problem

Input: Configuration C', traffic scenario F^S and capacity uncertainty W^S

Initialize: Set $r'_f := \arg\min_{r \in R^f} d_r^f$ for $f \in F^S$, Lagrange Multiplier $\mu_l := 0$ for $l \in L'$

Establish feasible solution: Iterate until $\bar{k}_l \leq 1 \forall l \in L'$

  Compute relative "weight" $w_{frl} = \sum_{e \in E^l} b_{freu} / \kappa_l^S$ for $f \in F^S, r \in R^f, l \in L'$

  Compute relative capacity shortage $\bar{k}_l = \sum_{f \in F^S} w_{fr'l}$ and set $l^* := \arg\max_l \bar{k}_l$.

  For flights with $w_{fr'l^*} > w_{frl^*}$ on l^*, store $\gamma_r^f = \frac{d_r^f - d_{r'}^f - \sum_{l \in L'} \mu_l (w_{fr'l} - w_{frl})}{w_{fr'l^*} - w_{frl^*}}$ for $r \in R^f$.

  Determine flight and route with lowest γ_r^f, update $r'_f = r$ and $\mu_{alu^*} = \mu_{alu^*} + \gamma_r^f$

Improve feasible solution: Iterate until no further improvement found, i.e., $\Delta d = \emptyset$

  For flights and routes with $d_{r'}^f > d_r^f$ and $\bar{k}_l - w_{fr'l} + w_{frl} \leq 1, l \in L'$, store $\Delta_r d = d_{r'}^f - d_r^f$.

  Find flight and route with largest $\Delta_r d$ and update $r'_f := r$.

Output: Routing $R^* = \{r'_f : f \in F^S\}$ and displacement cost $D^* = \sum_{r \in R^*} d_r$

3.C Case study data

Table 3.C.1. Capacity-side network characteristics in the case study

Airspace	Elementary sectors	Collapsed sectors	Configurations	Probability of ATFM regulation	Technology provider
EDUUUTAC	11	14	13	10.8%	Indra
EDUUUTAE	10	14	13	8.3%	Indra
EDUUUTAS	12	29	13	41.6%	Indra
EDUUUTAW	10	12	11	2.5%	Indra
EDYYBUTA	8	13	10	0.0%	Indra
EDYYDUTA	9	12	7	0.0%	Indra
EDYYHUTA	12	19	12	0.0%	Indra
EPWWCTA	18	77	26	56.8%	Indra
LHCCCTA	10	24	7	0.0%	Thales
LKAACTA	6	9	6	0.2%	Legacy
LKAAUTA	6	9	8	0.0%	Legacy
LOVVCTA	26	58	21	1.2%	Thales
LSAGUTA	6	9	11	3.1%	Legacy
LSAZUTA	6	7	7	0.3%	Legacy
LZBBCTA	27	69	8	0.0%	Thales

Technology providers taken from Eurocontrol (2022); "legacy" implies that ANSP employs its own system.

Table 3.C.2. Aircraft size matching based on ICAO classification.

Aircraft type	**Aircraft size**	**ICAO Wake Turbulence Category**
B733	medium	M (medium)
B738	medium	M (medium)
B75	medium	M (medium)
B76	large	H (heavy)
B74	large	H (heavy)
A319	medium	M (medium)
A320	medium	M (medium)
A321	medium	M (medium)
AT4	medium	M (medium)
AT7	medium	M (medium)
DH8	medium	M (medium)
E179	medium	M (medium)
A33	large	H (heavy)
A38	large	H (heavy)
A34	large	H (heavy)
B77	large	H (heavy)
B78	large	H (heavy)
BE20	small	L (light)
C560	small	L (light)
SMP	small	L (light)
EXMT	small	L (light)

SMP refers to other small aircraft types, and EXMT refers to flights typically exempt from demand management measures.

Table 3.C.3. Geographic setup for capacity sharing settings.

Cross-ACC

Alliance	Airspaces		
1	EDUUUTAC (G. Central)	EDUUUTAE (G. East)	
	EDUUUTAS (G. South)	EDUUUTAW (G. West)	
2	LKAACTA (Czech Rep.)	LKAAUTA (Czech Rep.)	
3	LSAGUTA (Switzerland)	LSAZUTA (Switzerland)	
4	EDYYBUTA (Maastricht)	EDYYDUTA (Maastricht)	EDYYHUTA (Maastricht)
5	LOVVCTA (Austria)		
6	LZBBCTA (Slovakia)		
7	EPWWCTA (Poland)		
8	LHCCCTA (Hungary)		

Cross-border

Alliance	Airspaces	
1	EDUUUTAE (G. East)	EDUUUTAS (G. South)
1	EDUUUTAW (G. West)	LKAAUTA (Czech Rep.)
2	EDYYBUTA (Maastricht)	LKAACTA (Czech Rep.)
3	EDYYDUTA (Maastricht)	EPWWCTA (Poland)
4	EDYYHUTA (Maastricht)	LOVVCTA (Austria)
5	LSAGUTA (Switzerland)	LHCCCTA (Hungary)
6	LSAZUTA (Switzerland)	LZBBCTA (Slovakia)
7	EDUUUTAC (G. Central)	

Common technology

Alliance	Airspaces	
1	EDUUUTAC (G. Central)	EDUUUTAW (G. West)
1	EDUUUTAE (G. East)	EDYYDUTA (Maastricht)
2	EDUUUTAS (G. South)	EDYYBUTA (Maastricht)
3	LKAACTA (Czech Rep.)	LKAAUTA (Czech Rep.)
4	EDYYHUTA (Maastricht)	EPWWCTA (Poland)
5	LZBBCTA (Slovakia)	LHCCCTA (Hungary)
6	LSAZUTA (Switzerland)	LSAGUTA (Switzerland)
7	LOVVCTA (Austria)	

Germany abbreviated by "G.".

Chapter 4

Environmental concerns: Demand-capacity balancing to reduce emissions

[1]

[1]Submitted in January 2022 to *Transportation Research Part A* under "Leveraging demand-capacity balancing to reduce air traffic emissions and improve overall network performance" with co-authors Arne K. Strauss, Nikola Ivanov, Radosav Jovanovic and Frank Fichert, see Künnen et al. (2022)

Abstract

Improvements in air traffic management (ATM) and aircraft operations may reduce European aviation emissions by 6% until 2050, compared to 2018 levels. Next to trajectory optimisation, in which more fuel-efficient trajectories are identified for one flight at a time, the improved balancing of demand and capacity in air traffic management can contribute to this reduction - a lever that has not received much research attention so far. Today, many flights in Europe are diverted from shorter trajectories due to insufficient capacities in the network. In addition, some airspace users (AUs) deliberately plan to fly on longer trajectories to minimise cost by avoiding countries with high en-route airspace charges. In both cases high cost are created for the network (in terms of delay and rerouting cost) and the environment.

In this paper, we analyse two demand and capacity management mechanisms that aim at improving flight efficiency and reducing emissions. On the capacity side, we evaluate the impact of centrally-determined capacities (instead of locally-determined capacities) that consider both network effects and emissions. On the demand side, we analyse the effect of trajectory-independent airport-pair charges (instead of country-specific airspace charges) that align the incentive of AUs with the environment when making their trajectory choice. The mechanisms are tested on a realistically-sized case study covering 3,000-4,000 flights in large parts of Western European airspace. We find that central capacity planning can reduce variable network cost by 21% and emissions from detours by almost 64%. Furthermore, airport-pair charging can save almost 11% of variable network cost and up to 320,000 tons of CO_2 emissions if accompanied by capacity changes that reflect the shift in demand towards shorter trajectories.

Keywords: air traffic management; capacity planning; demand-capacity balancing; aviation emissions

4.1 Introduction

Air traffic is one of the largest sources of greenhouse gas emissions, accounting for around 2.4% of all antrophogenic CO_2 emissions in 2018 (Lee et al. (2021)). As such, governments and international organizations alike have called for measures to reduce emissions from air traffic. In Europe, where air traffic accounts for almost 4% of total emissions (European Commission (2021)), ambitious targets for emission reduction in the transport sector have been set. The European Green Deal requires that transport emissions are reduced by 90% until 2050, based on 1990 emissions (European Commission (2021)). At the same time, demand for air traffic is expected to grow further in the next decades. Aviation experts suggest that global passenger numbers may triple by 2050, even after accounting for the impact of the Covid-19 pandemic (International Civil Aviation Organization (2021)).

Assuming a given number of flight movements, the technological options to reduce aviation emissions are limited: i) using more energy-efficient aircraft, ii) substituting kerosene with sustainable aviation fuels, or iii) improving air traffic management (ATM) and aircraft operations. Destination2050 estimates that ATM improvements can save up to 6% of yearly CO_2 emissions in Europe until 2050 (Destination2050 (2021)). As Matthes et al. (2020) noted, "a more sustainable ATM needs to integrate comprehensive environmental impacts and associated forecast uncertainties into route optimisation in order to identify robust eco-efficient trajectories". Furthermore, Baumgartner et al. (2021) suggest in a recent report that improved air navigation services are key to reduce aviation emissions, particularly in the short and medium term. That is why the Single European Sky ATM Reseach (SESAR) initiative in Europe has moved towards broad environmental impact reduction objectives including noise, air quality and climate change, next to existing measures on fuel and delays. In particular, SESAR has included an initiative to "minimise the environmental footprint of aviation" through improved flight trajectories as part of their Master Plan 2020 (SESAR Joint Undertaking (2020)).

An emission reduction measure that has so far not received much research attention falls under the category "improving air traffic management", namely, the balancing of capacity and demand for air traffic. Failing to provide capacity where and when needed could lead to demand shifts to longer routes (hence, more emissions) to avoid long delays in congested parts of the network. Eurocontrol estimates that fuel-inefficient routing across all phases of the flight was responsible for 8.6%–11.2% of CO_2 emissions from European air traffic in 2019 (Eurocontrol (2020c)).

The process of establishing demand-capacity balance (DCB) in Europe involves three key

stakeholders: Airspace Users (AUs) who generate demand for ATM services, Air Navigation Service Providers (ANSPs) who provide capacity, and the Network Manager (NM) which acts as a coordinator in the European network. The DCB process itself spans three phases: strategic, pre-tactical and tactical. Capacity decisions for a day of operation are usually made several weeks or months in advance (strategic phase). In particular, each ANSP decides on the number of Air Traffic Control Officers (ATCOs) to deploy in the airspace that they control, which in turn determines how many flights can cross the airspace in each time interval. In some cases, the capacity decision can be adjusted several days ahead of the day of operations (pre-tactical phase) to meet the anticipated demand, but this largely depends on capacity providers and their resources and capabilities to increase capacity at a short notice (e.g., by calling in more ATCOs). Demand decisions, on the other hand, are typically made by AUs on the day of operations (tactical phase). By submitting their flight plans the AUs "reveal" to the NM and ANSPs where and when in the network they would fly, which often happens only several hours before departure. Based on the latest information on demand and capacity in the network (i.e., weather, military activity and available ATCOs), ANSPs decides how many and which Air Traffic Control (ATC) sectors to open (the so-called sector opening scheme). Furthermore, the NM and AUs decide together on the final trajectories of each flight. If a flight cannot be routed on its preferred trajectory due to capacity shortage in the network, the NM can propose to delay the flight on-ground. The AU can then either accept the delay or choose an alternative trajectory instead.

The current DCB process features two shortcomings which can result in higher en-route flight emissions. First, the capacity decisions are taken locally by the ANSPs, without considering network-wide conditions. The capacities are often insufficient to accomodate upcoming traffic, which sometimes results in AUs choosing to fly longer trajectories to avoid congested portions of airspace or imposed ground delays. Second, the demand for air traffic services is determined by AUs who make their trajectory choice accounting for en-route airspace charges, amongst other factors, which differ for each country in Europe. Since the differences in en-route charges even between neighbouring countries can be substantial, some AUs may deliberately choose to fly longer trajectories to avoid expensive airspaces, thereby increasing flight emissions according to Delgado (2015).

In this paper, we determine the impact of improved capacity and demand management of air traffic in Europe on en-route flight efficiency and emissions. On the capacity side, we evaluate the effect of strategic capacity planning which considers network effects in the capacity decision (instead of local ANSP decisions). The methodology to make such network-oriented

capacity decisions has been presented in Künnen et al. (2021). We develop a holistic emission cost function that reflects both CO_2 and non-CO_2 related flight emissions to include it in the decision making. On the demand side, we evaluate the impact of trajectory-independent airport-pair charges (instead of en-route airspace charges) on trajectory choice and eventually emissions. We split the analysis of demand and capacity mechanisms in two parts, in which either the NM or the AU holds the main decision making power over final trajectories. Our main contributions lie in a) incorporating CO_2 and non-CO_2 emissions in strategic capacity planning, b) assessing the effect of these capacity decisions, as well as of airport-pair charges, on cost and emissions in the European air traffic network, and c) deriving policy implications from the observed effects.

The paper is organized as follows: In §4.2 we review relevant literature on ATM mechanisms that reduce flight emissions. In §4.3 we discuss the proposed improvements to demand-capacity balancing in European ATM and the methodology with which these improvements are evaluated. In §4.4 we present the case study for our analysis. We present the findings under NM decision-making in §4.5 and under AU decision-making in §4.6, and close with policy implications and recommendations for future research in §4.7.

4.2 Literature Review

With growing concerns on climate change in the last decades, research on the environmental impact of aviation has also picked up. While many researchers have addressed air quality and noise as environmental considerations in aviation since the 1980s, flight emissions have only received attention in the last 15–20 years. Williams et al. (2007) are one of the first to discuss how innovations in ATM can help mitigate the climate impact of aviation. The authors argue that total flight emissions will rise increasingly quickly as traffic increases if congestion and airborne delays cannot be reduced successfully. In particular, they suggest that a coordinated strategic planning of flights across Europe as well as improved demand-capacity balancing in the (pre-)tactical phase can improve air traffic flows. In the past years, two research streams in the area of sustainable ATM have developed that aim at improving flight efficiency and thus reduce emissions, namely trajectory optimization and route charging. In the following, we use the term route to denote an airport-pair, and trajectory to denote one (of potentially multiple) three-dimensional flight paths that an AU can choose for the route.

Trajectory optimization

Trajectory-based optimization (or TBO) is concerned with finding the "best" trajectory with regard to some performance indicator (e.g., delay or fuel cost), usually for one flight at a time. Hammad et al. (2020) give a broad review of articles that focus specifically on improving the sustainability of trajectories. In an early work, Clarke et al. (2008) discuss the trade-off between delays, fuel burn and emissions in en-route trajectory optimization. The authors analyse the impact of changes in air speed and heading angles to minimize fuel and delay cost. In this approach, emissions are included as a by-product of fuel burn but are not explicitly considered in the optimization. Hamdan et al. (2020) also investigate the trade-off between emissions and delays when making routing decisions. The authors develop a bi-objective integer program that simulatenously minimises CO_2 emissions and delay cost and test it on an illustrative numerical example.

More recently, researchers have developed a holistic approach to compute flight emissions (including non-CO_2 emissions and contrails) and incorporated the resulting cost function in the optimization to determine so-called climate-optimal trajectories. Rosenow et al. (2017) develop a trajectory optimization model that includes climate effects from both CO_2 and NO_x emissions by sequentially applying lateral and vertical trajectory optimization. Rosenow and Fricke (2019) use this emission-adjusted trajectory optimization model to test the impact of free routing conditions in Europe on network efficiency, safety and capacity. Matthes et al. (2017) further include contrail and cirrus formations to develop a multi-dimensional environmental change function with which trajectories are optimized. They find on a specific city-pair that if non-CO_2 effects are considered, the total climate impact can be reduced by more than 40% while fuel burn increases by 0.5%. Matthes et al. (2020) go beyond the effect of emissions and incorporate air quality and noise levels for a multi-criteria environmental assessment of trajectories. They test the concept on a one-day traffic sample in Europe. Finally, Grewe et al. (2014) determine trajectories which explicitly avoid climate-sensitive regions, i.e., regions in which the climate impact of non-CO_2 emissions is particularly large. The authors develop climate cost functions based on trajectories in the North Atlantic region, but do not comment on the emission reduction potential.

There are, however, two important differences between TBO and our approach: First, TBO focuses on the steering of flights in the airborne phase while we address the planning process in the three earlier phases (i.e., the strategic, pre-tactical and tactical phase of ATM). Second, the research on TBO aims at optimizing aircraft trajectories (either individually or an entire network) without considering capacity constraints. In contrast, we include airspace

capacity levels as key constraints in all analyses and even optimize over these levels. It is important to note that in order to analyse an entire European network structure, we take as given a set of operational trajectories between any two cities and optimize over the choice of such discrete trajectories. This is in contrast to research on trajectory optimization where the operational trajectories itself are subject of the optimization.

Route charging

Gillingwater et al. (2009) find that in some cases airlines decide to fly longer trajectories in order to save cost. Next to the avoidance of congested airspace or adverse weather conditions, airspace charges also play a role in these decisions. The authors analyse a total of 97 flight plans covering 14 routes in European airspace, and identify several trajectories with similar or slightly lower airspace charges but with longer distances, resulting in higher fuel burn and hence CO_2 emissions. Similarly, Delgado (2015) finds that around 6.4% of flights (in a sample of more than 10,000 flights) choose to fly trajectories which are more than five nautical miles longer than the shortest trajectory and have lower airspace charges. To address this issue, researchers and the industry have proposed various solutions. For instance, the Wise Persons Group (2019) argue that a unified route charge may reduce these disincentives and the respective environmental burden. Verbeek and Visser (2016) propose a new charging method called FRIDAY (Fixed Rate Incorporating Dynamic Allocation for optimal Yield), where route charges are calculated using the section lengths along the great circle distance. FRIDAY would effectively fix the route charge for a given airport-pair and aircraft type, which would remove the incentive to take detours.

To the best of our knowledge, no research has analysed the impact of such route charging schemes on detours and flight emissions. In this paper, we use the route charging principle first proposed in the COCTA research (COCTA Consortium (2017)), which applies a similar trajectory-independent charging principle as the FRIDAY method. However, it includes a mechanism to recover full capacity cost from a network rather than local perspective (for details, see Pavlović and Fichert (2019)). In contrast to existing research we analyse the impact of this charging scheme on cost (including delay and rerouting cost) and emissions in a realistic-sized case study.

Table 4.3.1. Description of analysed settings along DCB process.

	Capacity management (strategic)	**Demand management (pre-/tactical)**	
Setting		Charging principle	Routing decision
Baseline	Local capacity decision by ANSPs	Country-specific airspace charges	NM imposes delays, AU can choose re-routing instead
NM autonomy	Central capacity decision by NM	Trajectory-independent airport-pair charges	NM can impose delays or reroutings
AU autonomy	Local capacity decision by ANSPs	Trajectory-independent airport-pair charges	NM can only impose delays

4.3 Proposed demand-capacity balancing in Europe

With demand for air traffic services often exceeding en-route capacity in Europe (esp. in the peak summer months), effective demand-capacity balancing is particularly important in European airspace. In this section, we describe how the decision-making between the Network Manager (NM), airspace users (AUs) and air navigation service providers (ANSPs) can be structurally improved across the three phases of the DCB process to improve en-route flight efficiency and reduce emissions.

The current decision-making process and the proposed improvements are summarized in Table 4.3.1. As shown, the NM takes on a rather soft mediation role in the current system, since it has limited instruments available to influence capacity and demand decisions made by ANSPs and AUs. The local ANSPs decide on airspace capacities without considering demand in other parts of the network, which sometimes results in reroutings (and higher emissions) to avoid delays. Similarly, the airspace charges in each country are determined without considering adverse effects on detours and emissions. In addition, environmental considerations and in particular flight emissions do not currently form part of either capacity decisions or airspace charges.

To improve en-route flight efficiency in Europe and reduce emissions, we analyse two demand-capacity balancing mechanisms: network-oriented capacity decisions and trajectory-independent airport-pair charging. Which mechanisms is proposed depends on whether the NM or the AUs hold the main decision making power to select the final trajectory. In particular, if the nm decided on potential delays and reroutings of flights, the use of network-

oriented capacity levels may reduce capacity shortages and thus the need for the NM to reroute flights. In case the AU decides on the final trajectory, the use of airport-pair charges may reduce the incentive of AUs to choose longer trajectories in order to avoid countries with high en-route charges. The DCB processes for both of these settings, which we denote as *NM autonomy* and *AU autonomy* respectively, are detailed in the following.

In the setting *NM autonomy* we assume that the NM coordinates the capacity and demand management actions across European airspace. The concept was first proposed in COCTA Consortium (2017), and recently also advocated by Andribet et al. (2022) who propose a stronger role for the NM in order to "reinvent" European ATM after Covid-19. In this setting the NM decides how much capacity to order from each ANSP (considering network effects) in the strategic phase, and how to resolve any demand-capacity imbalance by either delaying or rerouting flights. We assume that AUs would choose the cost-minimising flight option in making their trajectory choices which, given the trajectory-independent airport-pair charges, is the shortest trajectory. The airport-pair charges would still serve to recover the cost of capacity provision. If the NM decides to reroute or delay a flight to resolve a capacity shortage, the AU pays the same airport-pair charges but incurs additional cost from fuel burn or delays. To analyse this setting, we need to determine the network-oriented capacities that the NM should order, and the routing of flights through the network such that the capacity limits are not exceeded. Even though the setting constitutes a rather large divergence from the current DCB process in Europe, analysing its impact on flight efficiency and emissions can give an indication of the improvement potential in the system. We also evaluate in this setting if sharing of capacities across Area Control Centers (ACCs) can help to reduce flight emissions even further. The idea is that ATCOs in one ACC are also eligible to control flights in one or more (neighbouring) ACCs to better accommodate fluctuations in demand. Künnen et al. (2021) find that such capacity-sharing can in fact significantly reduce the cost from reroutings and delays in European airspace.

In the setting *AU autonomy* we assume that the AUs largely hold the decision making power around which trajectories to fly for each of their flights. The capacities are still determined locally by each ANSP, and the AUs choose their preferred trajectories in the (pre-)tactical phase. In order to acknowledge the trajectory choice made by AUs, the NM can only impose ground delays on flights (if required due to the given capacity limits) but cannot reroute them. However, instead of current country-specific airspace charges we employ trajectory-independent airport-pair based charges. The use of airport-pair charges aligns the incentives of AUs with those of the environment since the shortest trajectory (based on

distance) will now also be the cheapest (based on fuel cost and ATM charges). This will *a priori* improve horizontal flight efficiency, measured in terms of trajectory deviation from the shortest route. To analyse this setting, we need to determine the trajectory choices made by AUs based on airport-pair charges vs airspace charges, as well as the delays to impose on flights such that the locally-determined capacity limits are not exceeded. In the following two subsections, we present the methodology with which we determine the improvement potential of network-oriented capacities under the *NM autonomy* and of airport-pair charges under the *AU autonomy*, respectively.

4.3.1 Capacity and routing decisions under NM autonomy

Our ultimate aim is to evaluate the cost and emission reduction potential from improved demand-capacity balancing. To judge the performance of our models, we first establish a baseline that reflects the current DCB process. Under *NM autonomy* we compare the network-oriented capacity decisions against reference capacity levels based on actual 2016 data. In particular, we use the most frequently used capacity profile in August as a reference for the high capacity setting, in October for the medium traffic setting and in April for the low traffic setting (see Table 4.A.4 in Appendix 4.A). These capacity levels reflect the local capacities provided by ANSPs and as such serve as relevant baselines for our analysis.

To determine network-oriented capacity levels, we use the approach proposed by Künnen et al. (2021). Operationally, the problem is modeled by deciding on the capacity level for each airspace that minimizes the sum of capacity, delay and rerouting cost. In the following, we refer to the sum of delay and rerouting cost as displacement cost. We consider the problem $\min_{x\in\mathcal{X}} \mathbb{E}_S(G(x,S)) + c^T x$, where $x \in \mathcal{X}$ is a vector of capacity levels for each airspace (in sector-hours) out of the finite set of such levels $\mathcal{X}$, S is a random scenario that reflects a materialization of demand and capacity uncertainties (due to non-scheduled flights, employee absence or adverse weather), $G(x,S)$ is the sum of displacement cost for a certain scenario S under capacity decision x, and vector c represents the capacity cost (per sector-hour) of each airspace. Determining the displacement cost $G(x,S)$ for a capacity decision x and scenario S requires solving the integrated routing and sector opening problem (IRSOP); the corresponding mixed integer program is shown in Appendix 4.B. Note that it is impossible to determine $\mathbb{E}_S(G(x,S))$ exactly due to the potentially infinite set of scenarios S over which the expectation is evaluated. The procedure proposed by Künnen et al. (2021) to approximate the solution to the capacity planning problem is summarized in Appendix 4.B. In order to incorporate emission considerations in the capacity decision, we modify their approach by

adding emission cost to the displacement cost in the evaluation of $G(x, S)$.

We also analyse how sensitive the results are with regard to changes in traffic levels and the price of CO_2. To evaluate different traffic levels, we simply adjust the traffic scenarios S in $G(x, S)$ that are used to determine capacity levels. We use an average of 4,000 flights (scheduled and non-scheduled) to create traffic scenarios S for the high traffic setting, 3,500 flights for medium traffic and 3,000 flights for low traffic; the distribution is based on actual monthly traffic in 2018. To assess the impact of different prices for CO_2, we recalculate emission cost of each trajectory based on CO_2 prices of 0 EUR and 100 EUR per ton (the upper limit was chosen based on market expectations for the price of CO_2 in 2030). We then reestimate $G(x, S)$ based on the adjusted emission cost.

As a final analysis, we want to assess the value of capacity-sharing (across two or more neighbouring ACCs) in improving flight efficiency and emissions. For this purpose, we define seven alliances where we pool together two or three ACCs with similar cost level, geography and infrastructure (see Table 4.3.2). We then compute total capacities for each alliance by summing the capacity levels of each airspace it contains, which were determined through the procedure described above. To decide how much of the available capacity in an alliance should be assigned to each of its airspaces in a given traffic and capacity scenario S, we solve the mixed integer program proposed by Künnen et al. (2021), see XCILP in Appendix 4.B. To single out the effect on displacement and emission cost, we assume that no additional variable capacity cost (per sector-hour) are incurred if an ATCO works outside his or her "home" airspace. This is because the majority of cost to establish capacity-sharing are fixed rather than variable cost (e.g., cost for infrastructure and training of ATCOs).

To evaluate the performance of both network-oriented capacities and reference capacities, we test them under various traffic and capacity scenarios S. For this purpose, we develop 100 test scenarios (which we fix for all tests) and determine the "optimal" routing of flights in each scenario using the heuristic proposed by Künnen and Strauss (2022), see Algorithm 6 in Appendix 4.B. We then approximate the expected displacement and emission cost associated with a capacity decision as the average cost observed over these scenarios. It is important to note that while in the standard setting (without capacity-sharing) the capacity levels of each airspace remain fixed across all 100 test scenarios, in the setting with capacity-sharing the levels are adjusted depending on the traffic and capacity scenario (only the total alliance capacities remain fixed).

Table 4.3.2. Overview of alliances for capacity-sharing.

Alliance	Airspaces		
1	EDUUUTAC (G. Central)	EDUUUTAE (G. East)	
2	EDUUUTAS (G. South)	EDUUUTAW (G. West)	
3	EPWWCTA (Poland)	LZBBCTA (Slovakia)	
4	LOVVCTA (Austria)	LHCCCTA (Hungary)	
5	LKAACTA (Czech Rep.)	LKAAUTA (Czech Rep.)	
6	LSAGUTA (Switzerland)	LSAZUTA (Switzerland)	
7	EDYYBUTA (Maastricht)	EDYYDUTA (Maastricht)	EDYYHUTA (Maastricht)

Germany abbreviated as "G.".

4.3.2 Determining trajectory choices under AU autonomy

To model the local capacity decisions in the strategic phase of the demand-capacity balancing process under *AU autonomy*, we use the same reference capacity levels as described above (i.e., based on historic data). Overall, we assume that AUs make their trajectory choices primarily based on cost, that is, we assume that AUs always opt for the cost-minimising trajectory option for each flight. In the current system, AU's trajectory choices also depend on the specific en-route airspace charges for each country, next to other cost considerations such as fuel burn and delay cost. Therefore, under current airspace charges the "cheapest" trajectory (based on cost) and shortest trajectory (based on distance) do not necessarily coincide. With the proposed airport-pair charging principle, in which any trajectory option between two airports incurs the same charge, AUs no longer face a trade-off between higher costs due to longer routes and lower cost of en-route charges (see also Delgado (2015)). Therefore, we assume that AUs would, *ceteris paribus*, always choose to fly the shortest trajectory. Note that during the day of operations, wind can affect the AU's trajectory choice as well, such that the shortest trajectory may not be the most fuel-efficient one (Zillies et al. (2014)). However, as long as AUs choose the most fuel-efficient trajectory (given the wind conditions), the flight's fuel burn and hence its impact on the environment is minimised, and we do not need to consider any cost from detours. The process to determine both shortest and "cheapest" trajectories for each flight using Eurocontrol's Network Strategic Tool (NEST) is described in §4.4.1.

Finally, to determine the routings in the tactical phase, we again use the routing heuristic proposed in Künnen and Strauss (2022) and test the process on the same 100 test scenarios as above. To ensure that the chosen trajectories can only be delayed (and not rerouted) in

the final routing, we provide only delay options (and no alternative trajectories) as input to the routing heuristic. We test the trajectory choices made based on airspace vs airport-pair charges also in the case with no capacity restrictions to analyse the maximum potential emission savings through airport-pair charging.

4.4 Case study

4.4.1 Dataset description

We test the proposed mechanisms using a case study with real traffic data covering up to 4,300 flights across 15 Area Control Centers (or ACCs) in 8 European countries. The main source of the case study is Eurocontrol's Demand Data Repository (DDR2). On the capacity side, the dataset contains 173 configurations of the 177 elementary sectors across the 15 ACCs. On the demand side, the data includes all flights crossing these 15 ACCs on September 9, 2016 (the day with most traffic in Europe in 2016); for our analysis, we limit the time window to 9am – 3pm. In Europe, around 80% of flights are scheduled flights whose schedules are published for the season ahead, which means that the NM knows the airport pairs and departure time of these flights already in the strategic phase. However, for the 20% of non-scheduled flights this information is not available, which creates uncertainty with regards to the overall traffic level and the spatio-temporal distribution of flights. In our traffic sample, we have a total of 3,500 scheduled flights. To generate a pool of non-scheduled flights, we randomly select 2,500 flights from all remaining flights in the dataset (inside or outside the 6 hour time window). A snapshot of exemplary flights in the network, created using NEST, is shown in Figure 4.4.1.

As mentioned above, we rely on Eurocontrol's NEST to generate the shortest and "cheapest" trajectory for each flight (accounting for network route structure, rules and constraints). In particular, NEST first finds (up to 40) horizontal route options for each flight based on the actual filed flight plans, and then calculates the vertical profile, taking into account the actual route network and constraints in both horizontal and vertical planes. We then identify the shortest trajectory for each flight based on total distance (without considering wind conditions, however). In contrast, to determine the "cheapest" trajectory we take into account both en-route charges and the trajectory's detour vs the shortest trajectory. In addition, we use NEST to determine up to 10 different alternative trajectories (to serve as rerouting options) which differ from the shortest trajectory in either horizontal or vertical dimension. To complete the trajectory options, we add five delay options for the shortest trajectory of each flight, with departure delays of 10, 20 and 30 minutes respectively. Note that flights

can be either rerouted or delayed, not both.

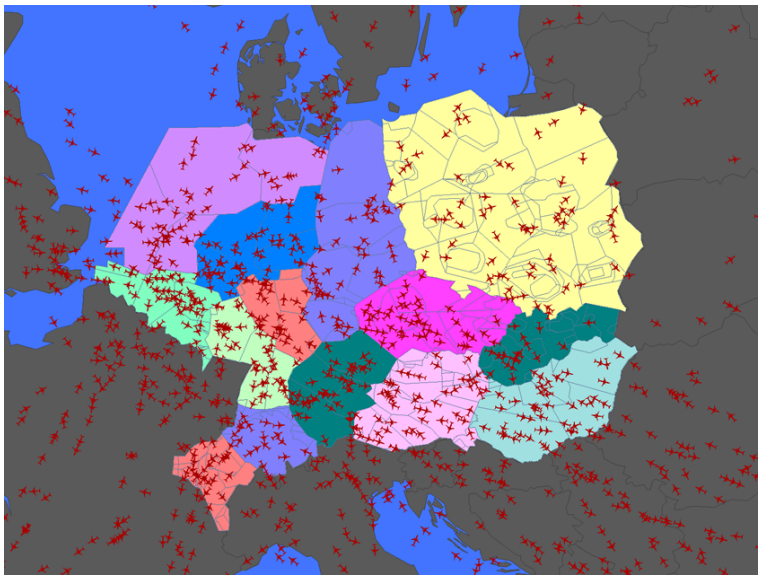

Figure 4.4.1: Snapshot of flights in the analysed airspace at a specific point in time.

Furthermore, we calculate cost for capacity, rerouting and delays. Capacity cost are computed as the average cost in each ANSP of opening one ATC sector for one hour, which are derived from annual cost reported by ANSPs. It is important to note that we only include ATCO cost in the case study (and treat them as variable) since they represent the capacity cost driver mostly affected by capacity decisions. Rerouting and delay cost are reported separately for each flight trajectory and differ by aircraft type. Rerouting cost are estimated based on Cook and Tanner (2015) and Eurocontrol (2018a); they include fuel, crew and passenger cost and increase linearly with the length of the detour (vs the shortest path). Delay cost are derived based on Cook and Tanner (2015) and increase non-linearly with the duration of delay. The emission cost associated with each trajectory is also estimated based on the length of the detour of the trajectory; the process is detailed in §4.4.2.

To reflect real-life uncertainties in air traffic, we include both demand and capacity variations. With around 20% of traffic coming from non-scheduled flights, demand cannot be known exactly in advance in the strategic phase (when capacity decisions are made). Therefore, we create various traffic scenarios which include, next to all scheduled flights, a random sample from the pool of non-scheduled flights. The number of non-scheduled flights selected from the pool in each scenario is drawn from a normal distribution with mean equal to 20% of total expected traffic and standard deviation of 100. Furthermore, we include capacity reductions due to employee absence (i.e., ATCO shortages) and adverse weather events. While

the former reduces the total number of available sector-hours in the impacted airspace, the latter reduces the actual capacity (i.e., how many flights can cross the sector during a certain time) of the impacted airspace sector. The probabilities for both types of events are derived based on historical data on ATFM regulations and are reported in the appendix. In case of an adverse weather in one of the airspaces, we assume that the capacity of a randomly selected elementary sector of that airspace (and the collapsed sector containing it), is reduced to 90% in 40% of cases, to 70% in another 40% of cases and to 50% in 20% of cases.

4.4.2 Measuring environmental impact

In order to make network-oriented capacity decisions that reflect full network cost, we develop a holistic emission cost estimation to include in our mechanism (next to capacity, delay and rerouting cost). We use a two-step approach to estimate emission cost for each trajectory in the dataset. First, we determine the emissions associated with a trajectory based on the flight distance vs the shortest trajectory. In our analysis we consider next to CO_2 also the impact of NO_x emissions, which represent by far the largest source of non-CO_2 emissions according to Lee et al. (2021). It has recently been shown that contrails and cirrus cloud formations can also have a sizable impact on global warming (see Bock and Burkhardt (2019)). However, we exclude them from our analysis because their occurrence and intensity cannot be reasonably predicted in the strategic stage of the DCB process (i.e., weeks to months in advance). Furthermore, we do not consider other environmental factors such as air quality and noise since their impact during the en-route phase of a flight is negligible. Both CO_2 and NO_x emissions increase approximately linearly with fuel burn. Lee et al. (2021) report conversion rates of 3.16 kg CO_2 and 15.1 g NO_x per kg of kerosene. The additional fuel consumption for each trajectory is computed based on the length of the detour vs the shortest trajectory (in nautical miles) and the aircraft's fuel consumption per nautical mile, which we take from Cook and Tanner (2015).

Second, we translate the emissions into a holistic cost function based on the climate impact of each emission source; the approach is proposed in Rosenow and Fricke (2019). NO_x emissions are converted into CO_2-equivalent emissions based on the Global Warming Potential of 20 years (or GWP20). The GWP20 measures the climate impact of greenhouse gases relative to that of CO_2 and reports a scaling factor for NO_x of 268 in Europe (see Lee et al. (2021)). Note that the impact of non-CO_2 emissions differs by the region in which they occur. We can then compute emission cost for each trajectory using total CO_2-equivalent emissions and the price of CO_2 based on the EU's emissions trading scheme (or

ETS), currently at around 65 EUR. Here, we assume that the ETS is extended to non-CO_2 emissions in the near future; this extension is already proposed by the European Union Aviation Safety Agency (EASA (2020)).

4.5 Demand-capacity balancing under Network Manager autonomy

In this section, we present the results from the proposed mechanism under *NM autonomy*. The performance of network-oriented capacities is evaluated based on both cost (i.e., capacity, displacement and emission cost) and emissions (i.e., CO_2 and NO_x emissions), and is compared against a relevant baseline which reflects the current local capacity decisions. We also include sensitivity analyses with regard to traffic levels and the price of CO_2.

Table 4.5.1 shows the cost performance of the network-oriented capacity levels (*Central*) against the identified baseline capacities (*Local*) on 100 scenarios. The capacity levels on which this assessment is based for both the *Central* model (determined using the methodology from §4.3.1) and the *Local* model are reported in Appendix 4.A. The models are evaluated on the high-traffic scenario with on average 4,000 flights across the 15 airspaces and 6-hour time window. The results show that the network-oriented capacities reduce variable network cost by over 26% compared to the baseline. Increasing variable capacity cost by only 2.7% can reduce the delay, rerouting and emission cost by around 54%. In particular, emission cost (from detours) are reduced by 44.6% which confirms the value of using network-oriented capacities to reduce environmental impact. We also find that in the baseline setting, emission cost amount to almost half of the other cost from detours (i.e., rerouting cost) and to around 10% of total variable cost, which is substantial considering that these cost are currently neglected in the DCB process. Recall that both rerouting and emission cost are calculated net of the shortest trajectory, i.e., we do not consider the fuel and emissions generated by flying the shortest trajectory itself. Figure 4.5.1 shows how the network-oriented capacities perform against the baseline (in terms of displacement and emission cost) for each of the 100 tested scenarios. While displacement and emission cost fluctuate between around 100,000 EUR and 250,000 EUR under baseline capacity levels, this range reduces to between 40,000 EUR and 120,000 EUR under network-oriented capacities.

With the cost improvements we also expect non-financial indicators to improve. Table 4.5.2 compares the delays and reroutings (incl. emissions) imposed on flights for the two models. We find that using network-oriented capacities reduces the share of delayed flights

Table 4.5.1. Cost performance of capacity models (high-traffic setting, n = 100 runs).

	Variable cost (EUR)				
Model	Capacity	Delay	Rerouting	Emission	Total
Local	154,287	69,304	60,724	32,311	316,626
Central	158,477	23,983	32,928	17,890	233,277
Difference	4,190	- 45,322	- 27,796	- 14,420	- 83,349
Difference (%)	+2.7%	-65.4%	-45.8%	- 44.6%	-26.3%

Rerouting and emission cost computed net of shortest trajectory.

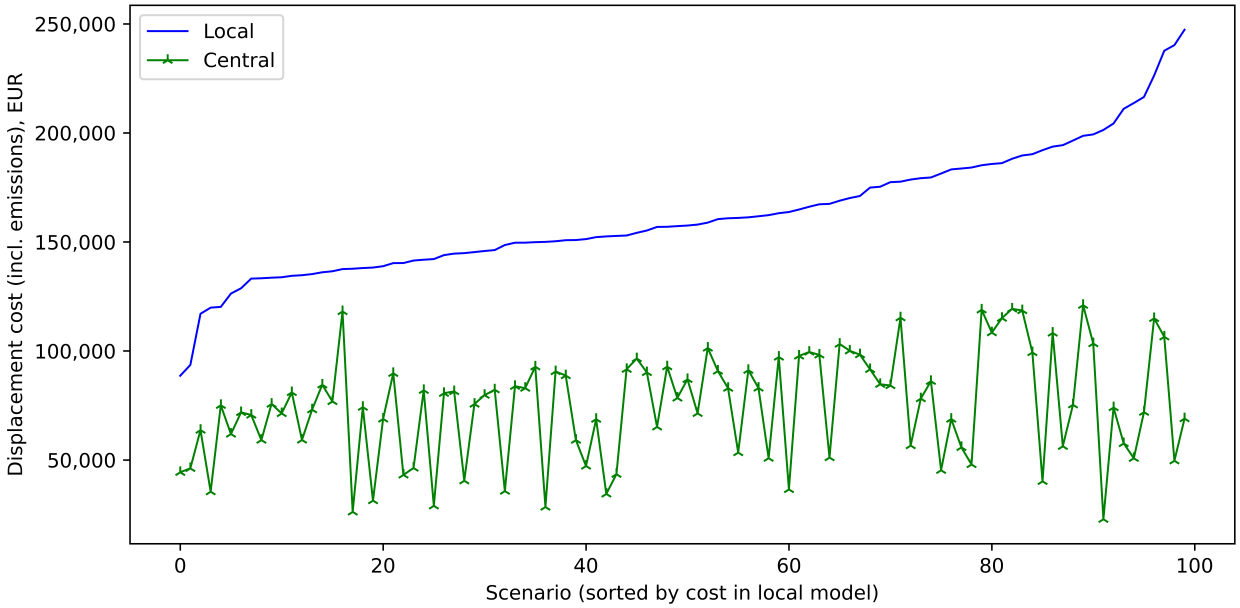

Figure 4.5.1: Displacement cost impact of capacity models across 100 scenarios.

from 4.7% to 2.5% and the average delay from almost 18 to under 15 minutes. Furthermore, the share of rerouted flights can be reduced from 21.6% to 16.8% and the average detour from over 13 to 9.2 nautical miles. This also implies that over 4 out of 5 flights can fly on the shortest trajectory (i.e., are neither delayed nor rerouted) with network-oriented capacities. Due to the lower amount of reroutings, CO_2 emissions can be reduced by almost 100 tons and NO_x emissions by around 0.5 tons. Both savings represent a reduction of around 45%; the relative changes of CO_2 and NO_x emission are equivalent since we assume both to increase linearly with fuel consumption. The fact that emissions are reduced by 45% while the average length of the detour is only reduced by 29% implies that particularly flights with above-average fuel consumption are not being rerouted anymore under network-oriented capacities.

Table 4.5.2. Overview of regulations for capacity models (high-traffic setting, n = 100 runs).

Regulation	**Delays**		**Reroutings**		**Emissions**	
Model/KPI	Share	Avg. delay	Share	Avg. distance	Total CO_2	Total NO_x
Local	4.7%	17.8 min.	21.6%	13.1 NM	217.6 t	1.04 t
Central	2.5%	14.6 min.	16.8%	9.2 NM	120.5 t	0.58 t
Difference (%)	-47.9%	-18.1%	-22.2%	-29.4%	-44.6%	-44.6%

Reroutings and emissions computed net of shortest trajectory.

To illustrate the purpose of identifying network-oriented capacities, we show in Figure 4.5.2 the trade-off between capacity and displacement cost (with and without emission cost) for an exemplary airspace: EDYYDUTA. While capacity cost are calculated directly from the capacity levels, the displacement and emission cost are estimated based on equation (4.5) in Appendix 4.B, and averaged across 100 scenarios. The graph shows that while capacity cost increase linearly with the capacity level (recall that we assume variable capacity cost), both displacement and emission cost decrease with capacities until a plateau is reached. The total cost curves show that the optimum point with lowest total cost is achieved for a capacity level of around 35 sector-hours, in the case with and without emission cost. In contrast, the baseline capacity level for EDYYDUTA for 2016 (which we identified based on historically used capacities in the peak summer period) was only 25 sector-hours. The analysis confirms that the historically used capacity levels are insufficient to cover high traffic volumes leading to many otherwise unnecessary reroutings and thus emissions.

4.5.1 Sensitivity with regards to traffic level and CO_2 price

To evaluate how sensitive the presented findings are with regards to the most important input parameters, we test the models for different traffic and CO_2 price levels. To analyse the impact of traffic intensity, we develop a low- and medium-traffic setting (next to the analysed high-traffic setting) with an average of 3,000 and 3,500 flights, respectively. As in the high-traffic setting, the actual number of flights differs in each scenario, but centers around the indicated average value. Figure 4.5.3 compares the cost performance of the capacity models in each of the three settings. The exact capacity levels and the full cost breakdown behind Figure 4.5.3 can be found in Appendix 4.A. We find that total variable network cost can be reduced substantially by using network-oriented capacities in all three settings, driven largely by a decrease in displacement and emissions cost. As expected, the potential decreases slightly for lower traffic levels, but the savings of around 50,500 EUR (or

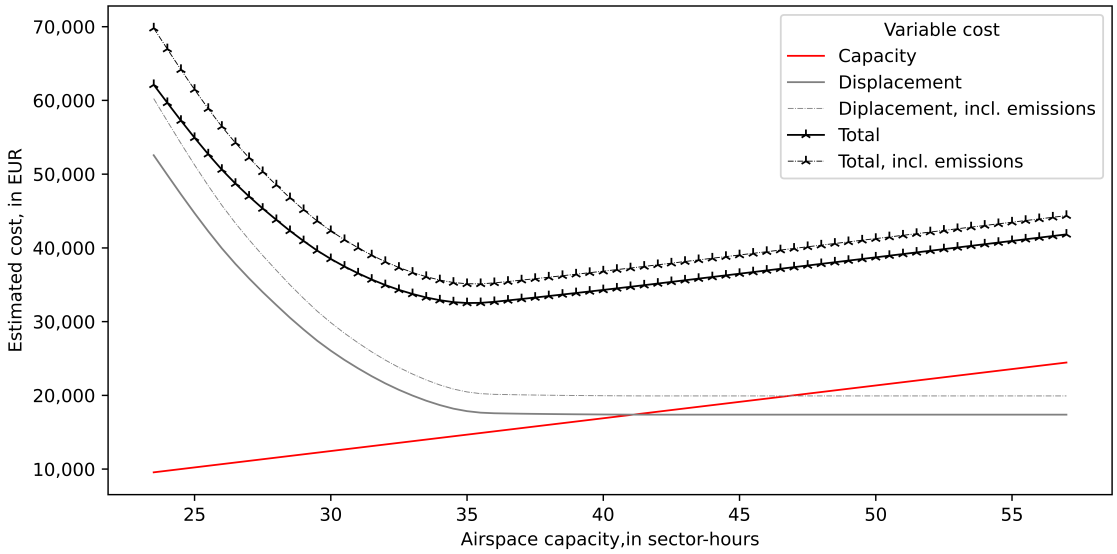

Figure 4.5.2: Trade-off between displacement and capacity cost (ex. for EDYYDUTA).
Displacement and emission cost are estimated using equation (4.5) in Appendix 4.B.

21.9% vs the baseline) under medium traffic and 27,900 EUR (or 15.9%) under low traffic are still sizable. If we assume that high-traffic levels occur for around three months of the year (i.e., peak summer period from June to August), low traffic for four months (November to February) and medium traffic for the remaining five months (based on historical traffic data from Eurocontrol (2019b)), we find that variable network cost can be reduced by 21% on average. Note that the displacement and emission cost for the *Central* model are already below 5,000 EUR for low traffic so that the improvement potential is limited for further decreases in traffic level.

In addition to financial performance, we are particularly interested in how well the proposed capacity model reduces CO_2 and non-CO_2 emissions. Table 4.5.3 shows the emission reduction across the three traffic levels. We find that CO_2 and NO_x emissions from detours can be reduced between 44.6% (under high traffic) and 82.6% (under low traffic). Assuming the same traffic levels distribution as before, we estimate that network-oriented capacities can save emissions from detours by 63.8% on average.

We also want to test how sensitive our results are with regard to the price of CO_2, which we use to compute overall emissions cost and thus to guide both our capacity and routing decisions. For that purpose, we develop a low price setting with no cost of CO_2 and a high price setting with 100 EUR per ton of CO_2 (next to the medium price of 65 EUR per ton used so far). The cost performance of the resulting network-oriented capacities is compared

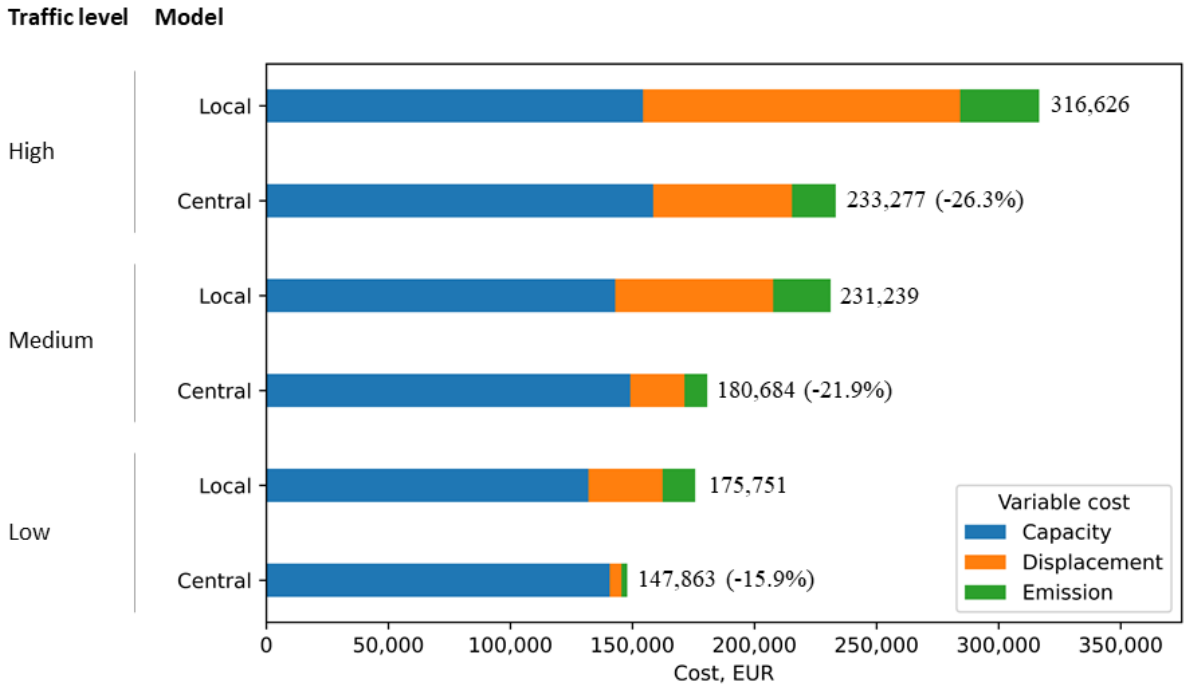

Figure 4.5.3: Cost performance of capacity models by traffic level.
Emission volumes and cost computed net of shortest trajectory.

against the baseline (assuming high traffic) in Figure 4.5.4. Again, the full cost breakdown can be found in the appendix. We find that, in contrast to traffic levels, the capacity decision is rather insensitive to changes in the price of CO_2. The variable capacity cost associated with network-oriented capacities increases only slightly from 156,600 EUR (for CO_2 price of 0 EUR) to 159,000 EUR (for a price of 100 EUR), while it changes quite substantially from 158,500 EUR (for high traffic) to only 140,700 EUR (for low traffic). This suggests that in order to make better capacity decisions it is more important to adequately predict upcoming traffic than to accurately estimate emission cost. However, Figure 4.5.4 also shows that the savings potential by using network-oriented capacities increases from around 19% (for CO_2 price of 0 EUR) to over 27% (for a price of 100 EUR). This suggests that making network-oriented capacity decisions becomes even more important with rising CO_2 prices.

4.5.2 Impact of capacity-sharing

In a final analysis, we want to test if sharing of capacities across ACCs in Europe can help to reduce total cost and emissions even further. Table 4.5.4 reports the cost performance of the

Table 4.5.3. Emissions savings by traffic level and type of emission (n = 100 runs).

Traffic level	**High**		**Medium**		**Low**	
Model/Emission	CO_2 (t)	NO_x (t)	CO_2 (t)	NO_x (t)	CO_2 (t)	NO_x (t)
Local	217.6	1.04	158.8	0.76	90.0	0.43
Central	120.5	0.58	63.1	0.30	15.7	0.08
Difference	-97.1	-0.47	-95.7	-0.46	-74.4	-0.36
Difference (%)	-44.6%	-44.6%	-60.2%	-60.2%	-82.6%	-82.6%

Emission volumes computed net of shortest trajectory.

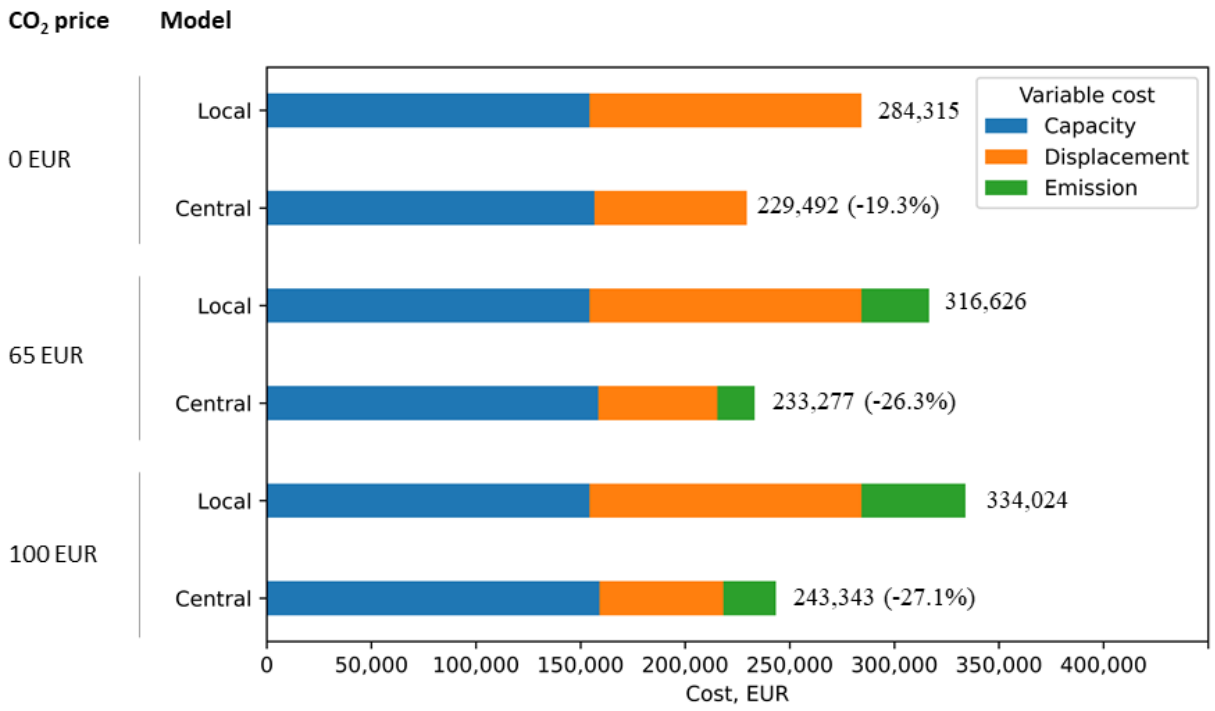

Figure 4.5.4: Cost performance of capacity models by CO_2 price level.
Emission volumes and cost computed net of shortest trajectory.

Central model with and without capacity-sharing in all traffic and CO_2 price settings. Recall that we allow capacity-sharing across the pre-defined alliances from Table 4.3.2 and that the total capacity for an alliance is the same as in the *Central* model without capacity-sharing. The results show that capacity-sharing can save between 2.0% and 5.1% of variable network cost depending on traffic levels (assuming a CO_2 price of 65 EUR). If the price of CO_2 increases to 100 EUR, the savings increase to 6.1%. Table 4.5.5 reports the corresponding emissions savings and shows that capacity-sharing can save an additional 10.1% – 38.5%

in CO_2 and NO_x emissions, compared with network-oriented capacities without capacity-sharing.

Table 4.5.4. Variable network cost with and without capacity-sharing (n = 100 runs).

	Traffic level			**Price of CO_2**	
Variable network cost	Low	Medium	High	0 EUR	100 EUR
Without capacity-sharing (EUR)	147,863	180,684	233,277	231,376	242,774
With capacity-sharing (EUR)	144,953	175,187	221,472	220,113	227,976
Difference (EUR)	-2,910	-5,497	- 11,806	-11,263	-14,799
Difference (%)	- 2.0%	- 3.0%	- 5.1%	- 4.9%	- 6.1%

Table 4.5.5. Emission savings of capacity-sharing by traffic level (n = 100 runs).

Traffic level	**High**		**Medium**		**Low**	
Emissions	CO_2 (t)	NO_x (t)	CO_2 (t)	NO_x (t)	CO_2 (t)	NO_x (t)
Without capacity-sharing	120.5	0.58	63.13	0.30	15.67	0.08
With capacity-sharing	108.4	0.52	52.27	0.25	9.65	0.05
Difference	- 12.1	- 0.1	- 10.9	- 0.1	- 6.0	- 0.0
Difference (%)	-10.1%	-10.1%	-17.2%	-17.2%	-38.5%	-38.5%

Emision volumes computed net of shortest trajectory.

Figure 4.5.5 shows the distribution of displacement cost (incl. emissions) with and without capacity-sharing for all three traffic settings. We find that the benefit from capacity-sharing increases with a) the traffic level in the network and b) the uncertainty in the underlying scenario. In particular, we find larger cost differences between the model with and without capacity-sharing under high traffic than under medium or low traffic, and for scenarios that show higher displacement cost (for all traffic settings). Furthermore, we observe a number of scenarios for each traffic setting in which capacity-sharing does not generate any savings. In these cases the actual demand and available capacity are in line with expectations so that the capacity distribution from the *Central* model does not need to be adjusted.

Rerouting and emission cost computed net of shortest trajectory.

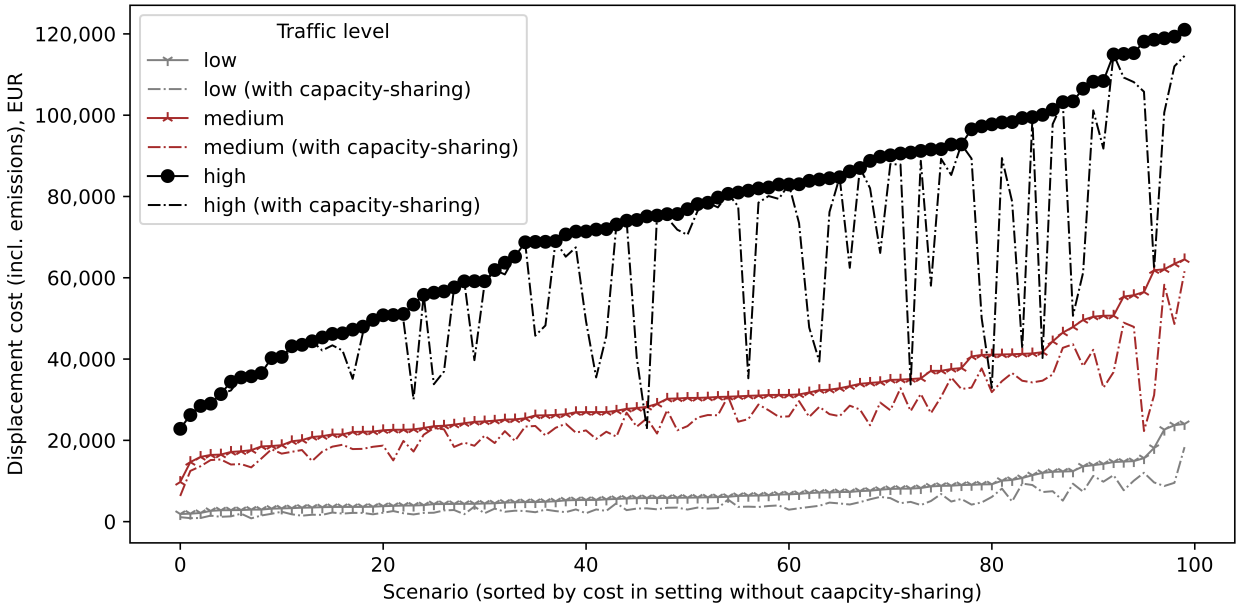

Figure 4.5.5: Displacement cost impact of capacity-sharing across 100 test scenarios.

Table 4.6.1. Key inefficiency metrics of trajectory choices under airspace charging.

Detours		**Emissions**	
Count	Avg. Distance	Avg. CO_2	Avg NO_x
5,109 (86%)	1.3 NM	31.7 kg	0.15 kg

4.6 Demand-capacity balancing under Airspace User autonomy

In this section, we analyse the impact of airport-pair charging on trajectory choices and eventually flight emissions, under AU autonomy. The impact is determined by comparing both cost and emissions against the trajectory choices made under current en-route airspace charging. Again, we test how sensitive our results are to changes in traffic levels and the price of CO_2. Table 4.6.1 summarizes the key inefficiency metrics of the identified trajectory choices based on current airspace charges (based on the methodology in §4.3.2) across all flights in the dataset. We find that 14% of flights choose to fly the shortest trajectory under airspace charging. All other flights choose a longer trajectory with an average detour of 1.3 nautical miles, creating additional emissions of 31.7 kg CO_2 and 0.15 kg NO_x per flight.

To evaluate the improvement potential through airport-pair charging, we use these trajectory choices made by AUs (under airspace charging) to determine total displacement and

emission cost for the different traffic level and CO_2 price settings. We then compare the cost performance with that under airport-pair charging, where we assume that AUs choose to fly the shortest trajectory. In a first analysis we assume that the network is not capacity-constrained, to determine the maximum improvement potential through airport-pair charging. In theory, this setting could be realized by increasing capacity levels such that demand can always be accommodated; in practice, the structural capacity limits of ANSPs may not allow such increases. The results in Table 4.6.2 show that total displacement and emission cost (from detours) from the trajectory choices under airspace charging vary between 34,769 EUR (for low traffic) and 45,593 EUR (for high traffic); these cost reduce to 28,647 EUR if CO_2 bears no cost, and almost double to 54,719 EUR if the price of CO_2 increases to 100 EUR (both for high traffic). Note that under airport-pair charging these displacement and emissions cost would effectively be reduced to 0 since without capacity constraints all flights can fly the shortest trajectory. Therefore, the figures in Table 4.6.2 represent the maximum potential cost savings from airport-pair charging.

Table 4.6.2. Cost under airspace charging, without capacity limits (n = 100 runs).

	Traffic level			**Price of CO_2**	
Cost	Low	Medium	High	0 EUR	100 EUR
Delay (EUR)	0	0	0	0	0
Rerouting (EUR)	21,831	25,473	28,647	28,647	28,647
Emission (EUR)	12,938	15,198	16,947	-	26,072
Total (EUR)	34,769	40,670	45,593	28,647	54,719

Rerouting and emission cost computed net of shortest trajectory.

Table 4.6.3 reports the corresponding emissions generated under airspace charging across the three traffic levels, which represent the maximum potential emission savings from implementing airport-pair charges. We find that the trajectory choices under airspace charging create between 87.2 and 114.2 tons of CO_2 and between 0.42 and 0.55 tons of NO_x. The comparison for different traffic levels also allows us to better estimate the average emissions savings potential from airport-pair charging across all intra-European flights. If we assume that our results scale from the 6-hour case study to all 35,000 daily flights within Europe under the high-traffic setting, to 30,000 flights under medium traffic and to 25,000 flights under low traffic (and assume the same traffic level distribution across the year as in §4.5), airport-pair charging could save up to 320,000 tons of CO_2 and 1,530 tons of NO_x per year.

In a second analysis we assume that the network is capacity-constrained. In particular,

Table 4.6.3. Emissions under airspace charging, without capacity limits (n = 100 runs).

	Traffic level		
Emission	Low	Medium	High
CO_2 *(t)*	87.2	102.4	114.2
NO_x *(t)*	0.42	0.49	0.55

Emission volumes computed net of shortest trajectory.

we set capacity levels based on the reference capacities of each traffic setting (see Appendix 4.A). In this case, the initial trajectory choices based on airport-pair or airspace charges may be delayed (but not rerouted, to acknowledge the trajectory choice) to satisfy capacity limits. Table 4.6.4 compares the cost performance of the resulting routings for the high traffic level. While no emission cost from detours are incurred under airport-pair charging (since all flights fly on the shortest trajectory), total variable network cost are actually 12% higher than under airspace charging. The increase is driven by large delay cost; in fact, flights are delayed by over 20 minutes on average. This suggests that the reference capacity levels are not suited for a situation in which all flights are choosing to fly the shortest trajectory. Figure 4.6.1 illustrates the shift from "cheapest" to shortest trajectory for three exemplary flights.

Table 4.6.4. Cost performance of demand models (high-traffic setting, n = 100 runs).

	Variable cost (EUR)				
	Capacity	Delay	Rerouting	Emission	Total
Airspace charges	154,287	386,767	28,644	16,945	586,643
Airport-pair charges	154,287	503,590	-	-	657,877
Difference	-	116,823	- 28,644	- 16,945	71,234
Difference (%)	-	+30.2%	-100.0%	-100.0%	+12.1%

Rerouting and emission cost computed net of shortest trajectory.

In order to reduce emissions without deteriorating network cost, airport-pair charging would have to be combined with adjusted capacity levels that reflect the shift in demand towards shorter trajectories. Therefore, we determine the required network-oriented capacity levels for airport-pair charging using the methodology in §4.3.1, and test the performance with adjusted capacity levels (see Appendix 4.A) against airspace charging. Figure 4.6.2 shows that total variable network cost can be reduced between 3% (for high traffic) and 16% (for low traffic), resulting in average savings of 10.8% under the assumed traffic level

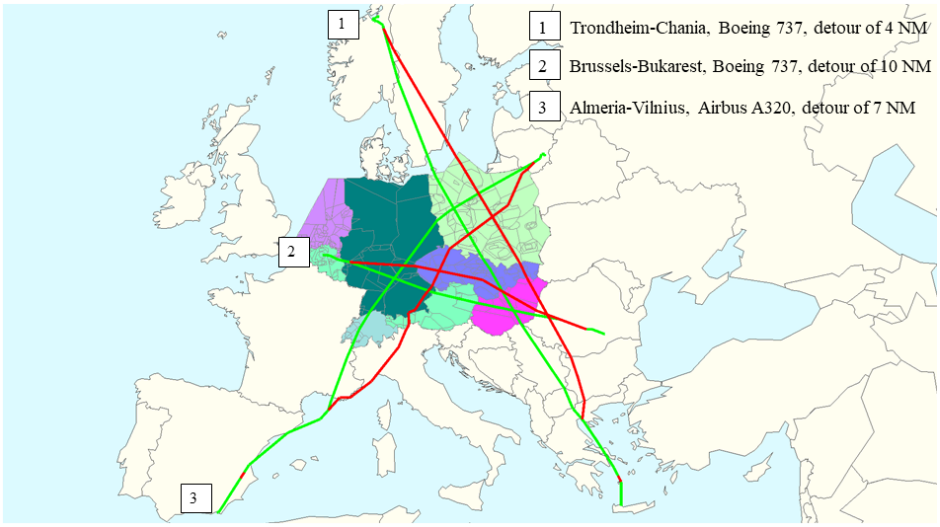

Figure 4.6.1: "Cheapest" (red) vs shortest (green) trajectory for three exemplary flights.

distribution.

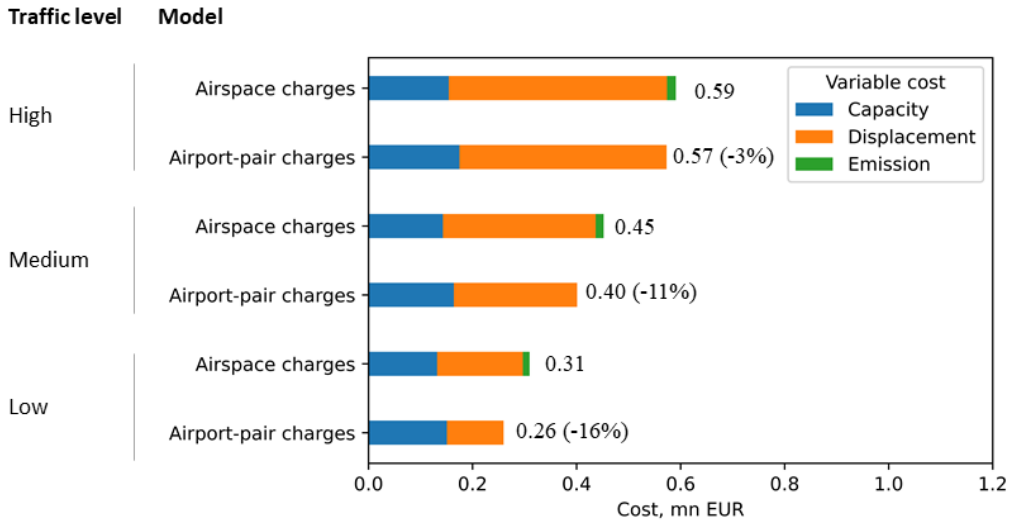

Figure 4.6.2: Performance of charging schemes after capacity adjustment (n = 100 runs).

4.7 Policy implications and conclusion

To guide decision makers towards improved en-route flight efficiency and reduced emissions in European ATM, we discuss in the following the policy implications and limitations of our results. Our findings suggest that making network-oriented capacity decisions could

substantially reduce total variable network cost. As exemplified for airspace EDYYDUTA in Figure 4.5.2, the current local capacity provisions are in some cases far from optimal, in terms of total variable network cost. In fact, every 1 EUR invested in additional capacity (provided in the right airspace at the right time) saves around 20.8 EUR in displacement and emission cost under the high traffic setting, around 9.2 EUR under medium traffic and still around 4.2 EUR under low traffic. Next to financial savings, network-oriented capacity decisions could reduce CO_2 and NO_x emissions from detours by almost 64% on average across the different traffic levels. To realize such savings, some central authority would need to coordinate capacity management across the European airspace. In one potential setup, such a central authority could be empowered to purchase capacities directly from ANSPs in line with expected demand in the network. Furthermore, our analysis shows that in order to make improved capacity decisions, it is more important to accurately predict upcoming traffic rather than to adequately estimate emission cost associated with reroutings in the network. Policy makers may also consider investing in infrastructure that allows for sharing of capacities (among neighboring ACCs or ANSPs). In our analysis, capacity-sharing further reduces total variable network cost by 2–5%, and emissions by 10–38% (depending on traffic level). However, while the savings potential is sizable, the legal and operational concerns to implement such capacity-sharing across the network are high and call for additional analyses (which go beyond the scope of this paper).

On the demand management side, policy makers may consider airport-pair charging as an effective but simple means to reduce emissions. We estimate that the trajectory choices made by airspace users based on current airspace charging create additional 320,000 tons of CO_2 and 1,530 tons of NO_x emissions per year. By aligning the incentives of airspace users and the environment, airport-pair charging may prevent environmentally-inefficient trajectory choices and save these emissions. The economic implications of this scheme are, however, harder to evaluate. While airport-pair charging can potentially save all rerouting and emission cost associated with these environmentally-inefficient trajectories, it also creates high delay cost under current capacity levels that offset these savings. Therefore, any implementation of airport-pair charges needs to be accompanied by an adjustment of capacity levels that reflects the shift in demand based on these charges. If capacity levels are adjusted accordingly, we find that total variable network cost can be reduced by almost 11% on average. It is important to note that the improvement potentials from the proposed capacity and demand management mechanisms do not substitute the achievements from trajectory-based optimisation, but rather complement it. The more (environmentally-)efficient the shortest

trajectory that trajectory-based optimisation can generate, the better the starting position for demand-capacity balancing in the strategic and (pre-)tactical phases of ATM.

There are a few limitations to our analysis that are important to consider in policy-making. First, the improvement potentials of both the capacity and demand management mechanisms depend on how well the baseline captures current decision making in the network. Second, for the potential of network-oriented capacities to be realised, we require that ANSPs are in fact able to increase their respective capacities (i.e., by hiring additional ATCOs). In some cases, the airspace design or the governance structure of an ANSP may not allow for the required increases. Third, we focus on capacity and demand management in the en-route part of flights. For the outlined potentials to be realised, the capacity and demand decisions in the terminal airspace and at airports would need to be coordinated accordingly. Finally, while we feel confident that the emission reduction potential that we observe in the case study scales to the network-wide level with 25,000–35,000 daily intra-European flights, such a linear scaling is less appropriate for variable network cost. This is because a) the network may be less congested outside the core European region that we analyse and outside the time window from 9am – 3pm (hence requiring less reroutings), and b) the cost of additional capacity may be higher than that of existing capacities (e.g., due to ATCOs working overtime). Future research may address the question of how cost improvements of both centralised capacities and airport-pair charges can be assessed for the entire European network.

Acknowledgements

This project has received funding from the SESAR Joint Undertaking within the framework of SESAR 2020 and the EU's Horizon 2020 research and innovation programme under the Grant Agreement Number 893380. The opinions expressed herein reflect the author's view only. Under no circumstances shall the SESAR Joint Undertaking be responsible for any use that may be made of the information contained herein.

Author Statement

Jan-Rasmus Künnen: Writing - Original Draft, Software, Methodology, Formal analysis, Investigation. **Arne Strauss**: Supervision, Conceptualization, Methodology, Writing – Review & Editing, Funding acquisition. **Nikola Ivanov**: Conceptualization, Data Curation, Writing – Review & Editing.**Radosav Jovanović**: Conceptualization, Funding acquisition, Writing – Review & Editing. **Frank Fichert**: Conceptualization, Funding acquisition.

Appendix

4.A Simulation results

Table 4.A.1. Performance of capacity models by traffic level (*NM autonomy*, n = 100 runs).

Traffic level	**High**		**Medium**		**Low**	
Cost/Model	Local	Central	Local	Central	Local	Central
Capacity	154,287	158,477	142,930	149,118	131,962	140,690
Displacement	130,029	56,910	64,730	22,193	30,423	4,846
Emission	32,311	17,890	23,579	9,373	13,366	2,327
Total cost	316,626	233,277	231,239	180,684	175,751	147,863

Table 4.A.2. Performance of capacity models by CO_2 price (*NM autonomy*, n = 100 runs).

Price of CO_2	**0 EUR**		**65 EUR**		**100 EUR**	
Cost/Model	Local	Central	Local	Central	Local	Central
Capacity	154,287	156,593	154,287	158,477	154,287	159,045
Displacement	130,029	72,899	130,029	56,910	130,029	59,267
Emission	-	-	32,311	17,890	49,709	25,030
Total cost	284,315	229,492	316,626	233,277	334,024	243,343

Table 4.A.3. Performance of demand models by traffic level (*AU autonomy*, n = 100 runs).

Traffic levels	**High**		**Medium**		**Low**	
Cost/Model	Airspace	Airport-pair	Airspace	Airport-pair	Airspace	Airport-pair
Capacity	154,287	154,287	142,930	142,930	131,962	131,962
Displacement	969,200	1,337,891	606,129	901,312	287,398	528,662
Emission	16,947	-	15,198	-	12,938	-
Total	1,140,433	1,492,177	764,256	1,044,242	432,299	660,623

Table 4.A.4. Overview of capacity levels by airspace and setting.

Setting	**Local (Baseline)**			**NM autonomy**			**AU autonomy**		
Airspace/Traffic	Low	Med.	High	Low	Med.	High	Low	Med.	High
EDUUUTAC	84	84	84	80	84	94	83	92	100
EDUUUTAE	72	72	73	72	72	75	77	87	98
EDUUUTAS	45	60	60	43	45	53	45	46	52
EDUUUTAW	72	72	72	66	69	77	66	72	79
EDYYBUTA	58	61	67	60	66	70	66	73	73
EDYYDUTA	45	47	50	50	60	65	54	63	66
EDYYHUTA	47	52	57	67	72	80	72	78	84
EPWWCTA	80	90	106	81	81	81	97	97	106
LHCCCTA	80	80	80	80	81	80	80	80	84
LKAACTA	50	52	50	44	44	44	44	44	44
LKAAUTA	36	42	56	36	36	41	36	36	36
LOVVCTA	99	104	114	101	100	100	112	122	136
LSAGUTA	32	40	47	43	50	54	50	57	59
LSAZUTA	34	42	50	48	50	54	52	57	59
LZBBCTA	39	48	56	39	46	43	39	46	44

4.B Simulation optimization approach to find network-oriented capacity levels

We consider the problem $\min_{x \in \mathcal{X}} \mathbb{E}_S(G(x,S)) + c^T x$, where $G(x,S)$ is the displacement cost for scenario S and capacity decision x, and vector c denotes the capacity cost (per sector-hour) of each airspace. The displacement cost $G(x,S)$ for a given capacity decision x and scenario S is then determined using the following mixed integer program:

$$G(x,S) = \min_{y,z} \sum_{f \in F^S} \sum_{r \in R^f} d_r^f y_r^f \tag{4.1}$$

$$\text{s.t.} \sum_{f \in F^S} \sum_{r \in R^f} b_{frlu} y_r^f z_{acu} \leq \kappa_l \qquad a \in A, c \in C^a, l \in L^c, u \in U \tag{4.2}$$

$$\sum_{c \in C^a} z_{acu} = 1 \qquad a \in A, u \in U \tag{4.3}$$

$$\sum_{r \in R^f} y_r^f = 1 \qquad f \in F \tag{4.4}$$

$$h_a \geq 0 \qquad \forall\, a \in A$$

$$y_r^f \in \{0,1\} \qquad f \in F, r \in R^f$$

$$z_{acu} \in \{0,1\} \qquad a \in A, c \in C^a, u \in U.$$

Given this problem definition, we determine the network-oriented capacity levels x^* in a two-step approach based on Künnen et al. (2021). In the first step, we determine linear regression coefficients that describe the relation between capacity shortage of each airspace (i.e., the number of flights that exceed the capacity limits of the airspace) and total displacement and emission cost in the network. For this purpose, an integer linear program (see CILP below) and a routing heuristic (see Algorithm 6 below) are used that approximate $G(x,S)$ for any combination of capacity level x and traffic and capacity scenario S. We run the heuristic for 2,000 different inputs for x and S (any lower number of inputs has shown to lead to unreliable results) and use the observed displacement cost $\hat{G}(x,S)$ to determine regression coefficients β based on equation (4.5). Since the running time of the routing heuristic increases quadratically with the number of available trajectory options for each flight (see Chapter 2), we pre-process the trajectories of all scheduled flights. In particular, we run the heuristic for 40 combinations of x and S and keep only those trajectory options which were chosen in the final routing decision at least once.

$$\hat{G}(x,S) = \beta_0 + \sum_{a \in A} \beta_a \sum_u k_a. \tag{4.5}$$

To ensure that the regression coefficients are not overfitted to the busy traffic patterns on the day of the case study but are also valid for lower traffic volumes, we uniformly sample between 3,000 and 4,000 total flights to create the traffic scenarios in S. This distribution

was chosen based on the actual monthly traffic distribution in the Eurocontrol area in 2018 (Eurocontrol (2019b)). To determine capacity shortages $k_a (a \in A)$ of each airspace, we assign each flight of the given scenario to its shortest trajectory, compute the resulting total traffic flows (in terms of number of flights) of the airspace and subtract from it the airspace capacity.

$$\textbf{(CILP)} \quad \min_{z} \sum_{a,c,u} k_{acu} z_{acu} \tag{4.6a}$$

$$\text{s.t.} \sum_{u} \sum_{c \in C^a} \bar{h}_{ac} z_{acu} \leq x_a^S \qquad a \in A \tag{4.6b}$$

$$\sum_{c \in C^a} z_{acu} = 1 \qquad a \in A, u \in U \tag{4.6c}$$

$$z_{acu} \in \{0,1\} \qquad a \in A, c \in C^a, u \in U. \tag{4.6d}$$

Algorithm 6 MMKP-based heuristic for routing problem

Input: Configuration C', traffic scenario F^S and capacity uncertainty W^S

1: **Initialize:** Set $r'_f := \arg\min_{r \in R^f} d_r^f$ for $f \in F^S$, Lagrange Multiplier $\mu_l := 0$ for $l \in L'$

2: **Establish feasible solution:** Iterate until $\bar{k}_l \leq 1 \forall l \in L'$

3: Compute relative "weight" $w_{frl} = \sum_{e \in E^l} b_{freu} / \kappa_l^S$ for $f \in F^S, r \in R^f, l \in L'$

4: Compute relative capacity shortage $\bar{k}_l = \sum_{f \in F^S} w_{fr'l}$ and set $l^* := \arg\max_l \bar{k}_l$.

5: For flights with $w_{fr'l^*} > w_{frl^*}$ on l^*, store $\gamma_r^f = \frac{d_r^f - d_{r'}^f - \sum_{l \in L'} \mu_l (w_{fr'l} - w_{frl})}{w_{fr'l^*} - w_{frl^*}}$ for $r \in R^f$.

6: Determine flight and route with lowest γ_r^f, update $r'_f = r$ and $\mu_{alu^*} = \mu_{alu^*} + \gamma_r^f$

7: **Improve feasible solution:** Iterate until no further improvement found, i.e., $\delta = \emptyset$

8: For flights and routes with $d_{r'}^f > d_r^f$ and $\bar{k}_l - w_{fr'l} + w_{frl} \leq 1, l \in L'$, store $\delta_r^f = d_{r'}^f - d_r^f$.

9: Find flight and route with largest δ_r^f and update $r'_f := r$.

Output: Routing $R^* = \{r'_f : f \in F^S\}$ and displacement cost $D^* = \sum_{r \in R^*} d_r$

In the second step of the capacity decision process, we use these regression coefficients in the exploration-exploitation algorithm proposed by Künnen et al. (2021), which is shown in Algorithm 7. In the procedure, a pool of promising capacity levels (or "candidates") is populated over time, and then iteratively reduced to the proposed best solution. In each iteration, the algorithm either tests an existing capacity levels from the pool on a new scenario to decide whether it remains in the pool, or it tests an entirely new capacity level to decide whether it is added to the pool. The procedure terminates if only one candidate remains in the pool, or if all candidates in the pool have already been tests on 300 different scenarios

S (in which case the best performing candidate is chosen). For a detailed description of the algorithm and notation, the reader is referred to Künnen et al. (2021).

Algorithm 7 Exploration-exploitation framework to seek optimum capacity

Input: Scenarios $\mathcal{S}$, solution space $\mathcal{X}$, sampling, resampling
Initialize: $i = 0$, $j = 1$, $\mathcal{X}_0^* = \emptyset$, $\hat{f}_0(x^*) = M$ (large number)
while $\exists x \in \mathcal{X}_j^* : N(x) < N^{max}$ **do**
  $i = i + 1$
  if $i = \lfloor j^{1.5} \rfloor$ **then**
    Select $x \in \mathcal{X}$ based on *sampling strategy*, evaluate $f(x, S_1)$ and set $j = j + 1$
    if $f(x, S_1) - \hat{f}_i(x^*) < \lambda$ **then**
      $\mathcal{X}_j^* = \mathcal{X}_{j-1}^* \cup \{x\}$, $\hat{f}_i(x) = f(x, S_1), N(x) = 1$
    else
      $\mathcal{X}_j^* = \mathcal{X}_{j-1}^*$
    end if
    for $x \in \mathcal{X}_j^*$ (ensure minimum amount of resampling) **do**
      if $N(x) < \lceil j^{0.5} \rceil$ **then**
        Set $n = \lceil j^{0.5} \rceil - N(x)$ and evaluate $f(x, S_{N(x)+1}), \ldots, f(X, S_{N(x)+n})$
        Set $\hat{f}_i(x) = \frac{\hat{f}_{i-1}(x)N(x) + \sum_{k=1}^{n} f(x, S_{N(x)+k})}{N(x)+n}$, and $N(x) = N(x) + 1$
      end if
    end for
    for $x \in \mathcal{X}_j^*$ (discard poor solutions) **do**
      if $\hat{f}_i(x) - \hat{f}_i(x^*) > \lambda / j^{0.15}$ **then**
        $\mathcal{X}_j^* = \mathcal{X}_j^* \setminus \{x\}$
      end if
    end for
  else
    Select $x \in \mathcal{X}_j^*$ based on *re-sampling strategy* and evaluate $f(x, S_{N(x)+1})$
    Set $\hat{f}_i(x) = \frac{\hat{f}_{i-1}(x)N(x) + f(X, S_{N(x)+1})}{N(x)+1}$, and $N(x) = N(x) + 1$
    Update $x_i^* = \arg\min_{x \in \mathcal{X}_j^*} \hat{f}_i(x)$ and
  end if
end while
Output: Capacity budget x_i^* with network cost $\hat{f}_i(x^*)$

$$
\begin{aligned}
\textbf{(XCILP)}\ \min_{h^0,z} & \sum_{a,c,u} \beta_a k_{acu} z_{acu} \\
\text{s.t. } & (4.6c), (4.6d) \\
& \sum_{u} \sum_{c \in C^a} \bar{h}_{ac} z_{acu} \leq x_a^S + h_a^0 && a \in A \\
& \sum_{a \in A^g} h_a^0 \leq |A^g| x_g^0 && g \in G \\
& h_a^0 \geq 0 && a \in A.
\end{aligned}
$$

4.C Case study data

Table 4.C.1. Capacity-side network characteristics in the case study.

Airspace	**Elementary sectors**	**Collapsed sectors**	**Configurations**	**Prob. of ATFM regulation**	
				Internal	External
EDUUUTAC	11	14	13	8.5%	5.5%
EDUUUTAE	10	14	13	1.4%	5.2%
EDUUUTAS	12	29	13	18.4%	8.5%
EDUUUTAW	10	12	11	1.1%	3.8%
EDYYBUTA	8	13	10	1.4%	6.6%
EDYYDUTA	9	12	7	0.8%	3.8%
EDYYHUTA	12	19	12	1.9%	3.3%
EPWWCTA	18	77	26	4.0%	1.9%
LHCCCTA	10	24	7	0.0%	3.8%
LKAACTA	6	9	6	0.0%	0.3%
LKAAUTA	6	9	8	0.0%	0.0%
LOVVCTA	26	58	21	1.4%	3.6%
LSAGUTA	6	9	11	1.4%	2.2%
LSAZUTA	6	7	7	0.5%	1.4%
LZBBCTA	27	69	8	0.0%	1.9%

ATFM regulations are divided into internal effects (e.g., employee absence) and external effects (e.g., adverse weather).

Chapter 5

Conclusions

Rethinking the demand-capacity balancing process in European air traffic management requires concepts for collaboration and coordination across the network. It remains an open (and political) question how the decision making power is best distributed among the key stakeholders in order to ensure efficient routings of flights. The concepts I present in this dissertation manage to overcome some of the challenges imposed on the network by fragmentation of airspaces (particularly in Europe) as well as the interplay between static capacities and volatile demand. I further develop the methodology to simulate and test different demand-capacity balancing mechanisms, such as flexible trajectory products, network-oriented capacity management and cross-border capacity sharing. Finally, I provide the mathematical tools that aid in making the complex decisions involved in capacity and demand management.

With this section I conclude the dissertation by summarizing the main contributions to both research and practice. Furthermore, I outline the limitations of the conducted research and induce areas for future research.

5.1 Contributions to research and practice

In Chapter 2 on demand management, the concept of dynamic pricing is extended to the field of ATM. The paper adds to existing research on ATM pricing by formulating the Dynamic Trajectory Pricing (DTP) problem, which is cast as a dynamic program. An efficient method is developed to determine the opportunity cost of so-called flexible trajectory products, based on which their dynamic prices are computed. The approach includes as a central component a MMKP-based routing heuristic to solve the integrated sector-opening and routing problem, which is shown to be $\mathcal{NP}$-hard. The heuristic reduces the solution time on a small instance

from 50 minutes to under 3 seconds, with around 11% accuracy compared to the exact approach. Furthermore, the two soft considerations included in the program to model revenue neutrality and fairness are not part of traditional revenue management problems and may prove valuable also outside the ATM context.

For practitioners, the chapter contributes a mathematical framework to evaluate different settings for the delegation of decision-making power between the network manager and airspace users to decide on flight trajectories. An algorithmic procedure is provided that determines the opportunity cost (and thus dynamic prices) of flexible trajectory products sufficiently fast, which makes the procedure suitable for real-time application. Finally, the small scale case study illustrates that flexible trajectory products provide an excellent compromise between the extremes of full decision making authority on parts of the network manager or airspace users: Implementing dynamically-priced flexible trajectory products achieves a network performance almost as if the network manager autonomously makes trajectory decisions while leaving the final trajectory choice to the airspace users. While today the routing of flights is adjusted rather ad-hoc in case of capacity shortages, flexible trajectory products help the network manager prevent capacity shortages in the first place by deliberately planning re-routings (and ground delays) for those flights that indicate a flexibility in their trajectories.

Chapter 3 on capacity management deals with the concept of capacity sharing as a mean to adjust capacities more flexibly to demand as well as uncertainties in capacity provision (e.g., due to adverse weather). Methodologically, the paper contributes to existing research a novel approach (based on the routing heuristic proposed in Chapter 2) to determine capacity levels for each airspace based on their stochastic performance. It is shown on a realistically-sized case study that the approach delivers higher-quality capacity decisions while requiring less computational resources than an existing deterministic procedure. Furthermore, the stochastic nature of the approach makes it suitable to analyze the value of capacity sharing as a hedge against an underprovision of capacity, which may have merit also outside the realm of ATM (e.g., in energy markets).

For practitioners and policy-makers in ATM, it remains an open question whether capacity decisions should be more centrally coordinated in order to reflect the network effects of such decisions. The findings in Chapter 3 show that such central, stochastically-determined capacities may substantially reduce variable network cost when compared to existing capacity levels. In addition, it may improve the resilience of the system by reducing the variation in network cost observed across different scenarios. The methodology introduced in the

paper also allows to evaluate different design options for capacity sharing, most notably the constitution of alliances for capacity sharing and provisions to prevent "outsourcing" of services to less costly providers. The case study shows that sharing capacities across borders (and service providers) achieves the largest reduction in variable network cost, followed by capacity sharing within the same service provider or of providers using the same technological infrastructure. Most importantly, these benefits can be reaped with only a small amount of cross-border resources, i.e., air traffic controllers that are trained to operate outside their local airspace. At the same time, the results vary strongly across different traffic intensities: if traffic volume is low, the likelihood of capacity underprovision and thus the benefit of capacity sharing is also low; and if traffic increases beyond a certain level, the value of capacity sharing decreases again because each airspace requires all their controllers themselves. The paper is the first to demonstrate the value of capacity sharing in a simulation study, and can provide guidance to ATM decision-makers on the merits, and the design, of such a scheme.

Environmental considerations are becoming increasingly relevant across sectors and regions, and ATM is no exception. Therefore, Chapter 4 discusses the effect of two demand-capacity balancing mechanisms on flight emissions. On the capacity side, flight emissions are explicitly included in the stochastic approach to capacity planning proposed in Chapter 3. On the demand side, the effect of trajectory-independent ATM charges on flight routing and emissions is analyzed. The paper provides substantial value to practitioners since it contributes a concept for making emission-adjusted capacity decisions and to analyze the environmental effects of capacity and demand decisions in ATM. A discussion with the Eurocontrol working group on Environmental Transparency in May 2022 highlighted that a structured approach to understand the interplay of capacities, cost and emissions was dearly needed. A case study on 3,000 flights across Western European airspace demonstrates that emission-adjusted capacities may substantially reduce the number of required flight detours (from their shortest trajectory) as well as the emissions associated with these detours. Finally, the existing airspace charging is estimated to create over 300,000 tons of additional CO_2 emissions in the European network per year, which a trajectory-independent charging could reduce. However, changing the ATM charging system would induce a geographic shift in demand in the network, which is why the distribution of capacities would need to be adjusted accordingly.

All in all, the dissertation shows that there is solid potential in rethinking the capacity and demand management decisions in European ATM, both conceptually and practically. After having revolutionized the revenue management of airlines in the past decades, dynamic

pricing may prove valuable also in the field of ATM. Linked with flexible trajectory products, the method can offer an effective compromise by improving network performance without fully shifting decision-making power to the network manager. Furthermore, a more centrally coordinated capacity planning that takes into account network effects and flight emissions can not only reduce total network cost, but may also improve the resilience of the system and reduce emissions. Capacity sharing within or across borders can intensify these benefits, and requires only a small amount of virtually deployed capacity.

5.2 Critical reflections and avenues for future research

The presented models contribute to existing research and demonstrate their practical value in numerical experiments. As in most research projects, the findings in this dissertation are subject to some limitations, which I would like to outline below. Many of these limitations point to future avenues of research, which I highlight accordingly.

With the pricing model proposed in Chapter 2, we demonstrate the value of flexible flight-to-route assignments through dynamically-priced trajectory products, which represents a novel approach for steering demand of air traffic services. The study features two main shortcomings: On the one hand, the pricing approach is tested on a rather small, artificial case study, and the proposed procedure does not scale adequately for realistic network sizes. As shown in Chapter 3 and 4, the routing heuristic which forms the basis of the pricing approach can in fact approximate the solution even to realistically-sized instances (i.e., more than 3,000 flights) in rather short time (i.e., less than one hour). However, the high frequency of such runs to regularly update opportunity cost estimates demands an even faster approach, at least if dynamic prices are to be computed in real-time. On the other hand, several key modeling assumptions and design choices are made in testing the pricing approach which limit the generalizability of results. Most importantly, a simplified airspace user choice model based on a binomial logit function is used to model the trajectory product decisions by airspace users when faced with different pricing offers. In addition, only two product types (one direct and one flexible product) are differentiated in testing the pricing approach, which limits both the potential and complexity of the procedure. Finally, the dynamic pricing approach is tested on only one scenario of capacity and demand intensity.

These limitations provide ample avenues for future research. First, researchers may build on the presented routing heuristic to develop more efficient approaches to price individual trajectories (or trajectory products) even in realistically-sized networks in real-time. Second,

a more nuanced version of the airspace user choice model would improve the reliability of numerical experiments on novel pricing approaches in ATM. Here, real-life observations or surveys may help understand how much airspace users are willing to pay for different levels of flexibility under various circumstances (with regards to e.g., the timing and criticality of a flight). Third, alternative pricing approaches (within both dynamic and differentiated pricing) and more complex product structures (covering more than two product types) may be investigated to capture the full potential of flexible trajectory products. Finally, to improve the robustness of results, the pricing approaches may be tested under various scenarios of capacity and demand intensity. This can help understand under which circumstances dynamically-priced trajectory products are most valuable.

In the following, I will jointly discuss the limitations to Chapters 3 and 4, since both chapters build on the same capacity planning model and case study. The limitations of the capacity planning model center around three core areas: scalability of the model, practical considerations, and scope of the analysis. First, while the proposed model is shown to deliver promising results on a large case study of more than 3,000 flights across Western Europe, it does not scale to a full day of operations across the European continent (with around 30,000 daily flights). This is because the regression parameters that are required as an input to the capacity model need to be estimated from at least 1,000 displacement cost observations. While not infeasible, generating these observations would require a prohibitive amount of computational resources. Furthermore, the capacity algorithm itself does not scale in the number of airspaces included in a capacity sharing alliance, which limits the number of alliance setups that can be investigated. Second, some of the modeling assumptions and design choices may not fully reflect practical considerations. Most importantly, the findings in Chapters 3 and 4 are derived based on capacity decisions for an isolated six-hour time frame (covering roughly one full ATCO shift). However, in practice, the capacity decision for one shift will affect the availability of controllers (and thus the potential capacity decision) in another shift, or even another day. Furthermore, it is assumed that each airspace can realize the determined, often increased capacity levels, but there may be practical limits as to whether these increased capacities may be sustained over time. Lastly, there are some practical obstacles that may impede implementation of capacity sharing in real-life. These include among others the regulatory environment in each country, licensing requirements of controllers, different air traffic technology equipment across airspaces, and contract design. Third, the scope of the analysis underlying Chapters 3 and 4 may be extended for future analyses. The presented capacity planning model is only applied to the en-route portion

of airspace, but the capacities of terminal airspace (particularly at slot-restricted airports) may strongly impact network performance as well. In addition, the environmental analysis in Chapter 4 does not consider the emissions impact of contrails, next to CO_2 and NO_x emissions, which may significantly alter the results.

Given these limitations, future researchers may further investigate the relationship between the level of capacity shortage in an airspace and the expected displacement cost (from delaying or rerouting flights accordingly). This may significantly improve the strategic ATM planning process since capacity shortages itself are rather simple and computationally inexpensive to predict. Furthermore, the presented capacity planning model may be adjusted to cover a longer time frame of up to one month or one year to enable truly long-term staffing decisions. Rostering practices that are specific to each airspace may be considered to make more practical recommendations. In order to further investigate the practical obstacles of the proposed capacity planning (and capacity sharing) models, the implementation may be piloted in a selected airspace. Finally, researchers may build on the presented models to develop a holistic capacity planning approach which takes into account the capacities in starting and terminal airspaces, alongside the considered en-route capacities.

In the past years, almost all industries have faced rising levels of uncertainty, driven by technological disruption, increasing environmental awareness and geopolitical tensions. For the aviation industry in particular, the sudden decline in air traffic demand during the Covid pandemic and the unexpected pace of recovery in 2022 have shown how sensitive the performance of the system has become to changes in demand. This experience highlights how important capacity and demand management measures are to ensure stable network operations in European ATM. The novel approaches to trajectory pricing and capacity planning presented in this dissertation may provide a first step to deliver on this promise.

Bibliography

S. Andradóttir and A. A. Prudius. Adaptive random search for continuous simulation optimization. *Naval Research Logistics (NRL)*, 57(6):583–604, Sept. 2010. doi: 10.1002/NAV.20422.

P. Andribet, M. Baumgartner, and J.-M. Garot. Reinventing european air traffic control based on the covid-19 pandemic experience. *Utilities Policy*, 75:101343, 2022. ISSN 0957-1787. doi: https://doi.org/10.1016/j.jup.2022.101343.

F. M. Baldursson, E. Lazarczyk, M. Ovaere, and S. Proost. Cross-border exchange and sharing of generation reserve capacity. *The Energy Journal*, 39(4), Oct. 2018. doi: 10.5547/01956574.39.4.FBAL.

C. Barnhart, D. Fearing, A. Odoni, and V. Vaze. Demand and capacity management in air transportation. *EURO Journal on Transportation and Logistics*, 1(1-2):135–155, June 2012. doi: 10.1007/s13676-012-0006-9.

M. Baumgartner, P. Andribet, and J.-M. Garot. Take advantage of the crisis to reinvent European air traffic control. Technical report, Union for Unity, Jan. 2021.

L. Bock and U. Burkhardt. Contrail cirrus radiative forcing for future air traffic. *Atmospheric Chemistry and Physics*, 19(12):8163–8174, 2019. ISSN 16807324. doi: 10.5194/acp-19-8163-2019.

T. Bolić, L. Castelli, and D. Rigonat. Peak-load pricing for the European air traffic management system using modulation of en-route charges. *European Journal of Transport and Infrastructure Research*, 17(1):136–152, 2017. ISSN 15677141. doi: 10.18757/ejtir.2017.17.1.3184.

L. Castelli, M. Labbé, and A. Violin. A network pricing formulation for the revenue maximization of European Air Navigation Service Providers. *Transportation Research*

Part C: Emerging Technologies, 33(June 2018):214–226, 2013. ISSN 0968090X. doi: 10.1016/j.trc.2012.04.013.

L. Castelli, T. Bolíc, S. Costanzo, D. Rigonat, E. Marcotte, and G. Tanner. Modulation of en-route charges to redistribute traffic in the European airspace. In *SESAR Innovation Days*, Bologna, 2015. Fifth SESAR Innovation Days.

J.-P. Clarke, M. Lowther, L. Ren, S. Solak, A. Vela, and L. Wong. En route traffic optimization to reduce environmental impact. *PARTNER Project 5 report*, (July), 2008.

COCTA Consortium. ATM value-chain redesign. Technical report, June 2017.

A. Cook and G. Tanner. European airline delay cost reference values. Technical Report December, University of Westminster, 2015.

L. Delgado. European route choice determinants: Examining fuel and route charge trade-offs. *Proceedings of the 11th USA/Europe Air Traffic Management Research and Development Seminar*, 2015.

Destination2050. Destination 2050 - a route to net zero European aviation. Technical report, Feb. 2021.

EASA. Updated analysis of the non-CO2 climate impacts of aviation and potential policy measures pursuant to the EU Emissions Trading System Directive Article 30(4). Technical report, 2020.

Eurocontrol. Performance Review Report (PRR 2017). Technical Report May, 2018.

Eurocontrol. Standard inputs for EUROCONTROL cost-benefit analyses. Technical report, Jan. 2018a.

Eurocontrol. ATM cost-effectiveness (ACE) 2016 benchmarking report with 2017-2021 outlook. Technical report, May 2018b.

Eurocontrol. Performance Review Report: An assessment of air traffic management in europe during the calendar year 2018. Technical report, 2019a.

Eurocontrol. Industry monitor 2018: the EUROCONTROL bulletin of air transport trends. Technical report, 2019b.

Eurocontrol. Performance Review Report: An assessment of air traffic management in europe during the calendar year 2019. Technical report, 2020a.

Eurocontrol. Technical note on en route capacity: a perspective on all aspects of en route capacity. Technical report, Dec. 2020b.

Eurocontrol. Environmental assessment: European ATM network fuel inefficiency study. Technical report, Dec. 2020c.

Eurocontrol. Local single sky implementation monitoring 2021. Technical report, May 2022.

European Commission. Legal, economic, and regulatory aspects of ATM data services provision and capacity on demand as part of the future European air space architecture, Dec. 2020.

European Commission. Reducing emissions from aviation. Technical report, 2021.

D. Gillingwater, L. C. Budd, R. E. Caves, and T. G. Reynolds. Environmental effects of aircraft operations and airspace charging regimes. final report. 2009.

V. Grewe, C. Frömming, S. Matthes, S. Brinkop, M. Ponater, S. Dietmüller, P. Jöckel, H. Garny, E. Tsati, K. Dahlmann, O. A. Søvde, J. Fuglestvedt, T. K. Berntsen, K. P. Shine, E. A. Irvine, T. Champougny, and P. Hullah. Aircraft routing with minimal climate impact: The REACT4C climate cost function modelling approach. *Geoscientific Model Development*, 7(1):175–201, 2014. ISSN 1991959X. doi: 10.5194/gmd-7-175-2014.

S. Hamdan, O. Jouini, and A. Cheaitou. Optimal air traffic flow management with carbon emissions considerations. *Springer Nature Switzerland*, 4:1078–1088, 2020. doi: 10.1007/978-3-030-21803-4.

A. W. Hammad, D. Rey, A. Bu-Qammaz, H. Grzybowska, and A. Akbarnezhad. Mathematical optimization in enhancing the sustainability of aircraft trajectory: A review. *International Journal of Sustainable Transportation*, 14(6):413–436, 2020. ISSN 15568334. doi: 10.1080/15568318.2019.1570403.

L. Hu and S. Andradóttir. An asymptotically optimal set approach for simulation optimization. *INFORMS Journal on Computing*, 31(1):21–39, Dec. 2019. doi: 10.1287/IJOC.2018.0811.

International Civil Aviation Organization. Post-COVID-19 forecasts scenarios. Technical report, 2021.

N. Ivanov, R. Jovanović, F. Fichert, A. Strauss, S. Starita, O. Babić, and G. Pavlović. Coordinated capacity and demand management in a redesigned Air Traffic Management value-chain. *Journal of Air Transport Management*, 75:139–152, Mar. 2019. doi: 10.1016/J.JAIRTRAMAN.2018.12.007.

R. Jovanović, V. Tošić, M. Čangalović, and M. Stanojević. Anticipatory modulation of air navigation charges to balance the use of airspace network capacities. *Transportation Research Part A: Policy and Practice*, 61:84–99, 2014. ISSN 09658564. doi: 10.1016/j.tra.2014.01.005.

R. Jovanović, O. Babić, and V. Tošić. Pricing to reconcile predictability, efficiency and equity in ATM. *Proceedings of the 11th USA/Europe Air Traffic Management Research and Development Seminar, ATM 2015*, 2015.

M. Khouja. The single-period (news-vendor) problem: literature review and suggestions for future research. *Omega*, 27(5):537–553, Oct. 1999. doi: 10.1016/S0305-0483(99)00017-1.

S. Kim, R. Pasupathy, and S. G. Henderson. A Guide to Sample Average Approximation. *International Series in Operations Research and Management Science*, 216:207–243, 2015. doi: 10.1007/978-1-4939-1384-8_8.

A. J. Kleywegt, A. Shapiro, and T. Homem-de Mello. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization*, 12(2):479–502, July 2002. doi: 10.1137/S1052623499363220.

J.-R. Künnen, A. Strauss, F. Fichert, N. Ivanov, R. Jovanovic, and S. Starita. Cross-border capacity planning in air traffic management under uncertainty. Unpublished working paper, Sept. 2021.

J.-R. Künnen, A. Strauss, F. Fichert, N. Ivanov, and R. Jovanovic. Leveraging demand-capacity balancing to reduce air traffic emissions and improve overall network performance. Unpublished working paper, Jan. 2022.

J.-R. Künnen and A. K. Strauss. The value of flexible flight-to-route assignments in pre-tactical air traffic management. *Transportation Research Part B: Methodological*, 160: 76–96, 2022. ISSN 0191-2615. doi: https://doi.org/10.1016/j.trb.2022.04.004.

D. S. Lee, D. W. Fahey, A. Skowron, M. R. Allen, U. Burkhardt, Q. Chen, S. J. Doherty, S. Freeman, P. M. Forster, J. Fuglestvedt, A. Gettelman, R. R. De León, L. L. Lim, M. T. Lund, R. J. Millar, B. Owen, J. E. Penner, G. Pitari, M. J. Prather, R. Sausen, and L. J. Wilcox. The contribution of global aviation to anthropogenic climate forcing for 2000 to 2018. *Atmospheric Environment*, 244(July 2020), 2021. ISSN 18732844. doi: 10.1016/j.atmosenv.2020.117834.

R. Levi, G. Perakis, and J. Uichanco. The data-driven newsvendor problem: new bounds and insights. *Operations Research*, 63(6):1294–1306, Oct. 2015. doi: 10.1287/OPRE.2015.1422.

S. Martello. Knapsack problems: algorithms and computer implementations. *Wiley-Interscience series in discrete mathematics and optimization*, 1990.

S. Matthes, V. Grewe, K. Dahlmann, C. Frömming, E. Irvine, L. Lim, F. Linke, B. Lührs, B. Owen, K. Shine, S. Stromatas, H. Yamashita, and F. Yin. A concept for multi-criteria environmental assessment of aircraft trajectories. *Aerospace*, 4(3):1–25, 2017. ISSN 22264310. doi: 10.3390/aerospace4030042.

S. Matthes, B. Lührs, K. Dahlmann, V. Grewe, F. Linke, F. Yin, E. Klingaman, and K. P. Shine. Climate-optimized trajectories and robust mitigation potential: Flying atm4e. *Aerospace*, 7(11):1–15, 2020. ISSN 22264310. doi: 10.3390/aerospace7110156.

M. Melgosa, X. Prats, Y. Xu, and L. Delgado. Enhanced demand and capacity balancing based on alternative trajectory options and traffic volume hotspot detection. In *2019 IEEE/AIAA 38th Digital Avionics Systems Conference (DASC)*. IEEE, 2019. doi: 10.1109/DASC43569.2019.9081803.

M. Moser, D. P. Jokanovic, and N. Shiratori. An algorithm for the multidimensional multiple-choice knapsack problem. *IEICE transactions on fundamentals of electronics, communications and computer sciences*, 80(3):582–589, 1997.

A. Mukherjee and M. Hansen. A dynamic rerouting model for air traffic flow management. *Transportation Research Part B: Methodological*, 43(1):159–171, 2009. doi: 10.1016/j.trb.2008.05.011.

M. Parlar, B. Rodrigues, and M. Sharafali. Event-based allocation of airline check-in counters: a simple dynamic optimization method supported by empirical data. *International Transactions in Operational Research*, 25(5):1553–1582, 2018. ISSN 14753995. doi: 10.1111/itor.12332.

G. Pavlović and F. Fichert. Effects of fragmentation on route charges. *Competition and Regulation in Network Industries*, 20:290 – 304, 2019.

D. Pisinger. A minimal algorithm for the multiple-choice knapsack problem. *European Journal of Operational Research*, 83(2):394–410, June 1995. doi: 10.1016/0377-2217(95)00015-I.

Y. Qin, R. Wang, A. J. Vakharia, Y. Chen, and M. M. Seref. The newsvendor problem: Review and directions for future research. *European Journal of Operational Research*, 213 (2):361–374, Sept. 2011. doi: 10.1016/j.ejor.2010.11.024.

T. Rambha and S. D. Boyles. Dynamic pricing in discrete time stochastic day-to-day route choice models. *Transportation Research Part B: Methodological*, 92:104–118, 2016.

A. Ranieri and L. Castelli. Pricing schemes based on air navigation service charges to reduce en-route ATFM delays. In *Third International Conference on Research in Air Transportation*, 2008.

J. Rosenow and H. Fricke. Impact of multi-criteria optimized trajectories on European airline efficiency, safety and airspace demand. *Journal of Air Transport Management*, 78(February):133–143, 2019. ISSN 09696997. doi: 10.1016/j.jairtraman.2019.01.001.

J. Rosenow, M. Lindner, and H. Fricke. Impact of climate costs on airline network and trajectory optimization: a parametric study. *CEAS Aeronautical Journal*, 8(2):371–384, 2017. ISSN 18695590. doi: 10.1007/s13272-017-0239-2.

SESAR Joint Undertaking. A proposal for the future architecture of the european airspace. Technical report, Feb. 2019.

SESAR Joint Undertaking. European ATM master plan executive view - digitalising Europe's aviation infrastructure. Technical report, 2020.

SESAR Joint Undertaking. Tests show operational feasibility of delegating air traffic services. Technical report, Apr. 2022.

S. Starita, A. K. Strauss, R. Jovanovic, N. Ivanov, and F. Fichert. Maximizing ATM cost-efficiency by flexible provision of airspace capacity. *SESAR Innovation Days*, 2016.

S. Starita, A. K. Strauss, X. Fei, R. Jovanović, N. Ivanov, G. Pavlović, and F. Fichert. Air Traffic Control Capacity Planning Under Demand and Capacity Provision Uncertainty. *Transportation Science*, 54(4):882–896, June 2020. doi: 10.1287/TRSC.2019.0962.

Steer Davies Gleave. Policy options for the modulation of charges in the Single European Sky. Technical Report April, 2015.

A. K. Strauss, R. Klein, and C. Steinhardt. A review of choice-based revenue management: Theory and methods. *European Journal of Operational Research*, 271(2):375–387, 2018. doi: 10.1016/j.ejor.2018.01.011.

R. J. D. Verbeek and H. G. Visser. Why aircraft will fly more fuel-efficiently on FRIDAY: The FRIDAY route charges method. *7th International Conference on Research in Air Transportation (ICRAT 2016)*, 2016.

V. Williams, R. B. Noland, A. Majumdar, R. Toumi, W. Ochieng, and J. Molloy. Reducing environmental impacts of aviation with innovative air traffic management technologies. *Aeronautical Journal*, 111(1125):741–749, 2007. ISSN 00019240.

Wise Persons Group. Report of the Wise Persons Group on the future of the Single European Sky, Apr. 2019.

Y. Xu, R. Dalmau, M. Melgosa, A. Montlaur, and X. Prats. A framework for collaborative air traffic flow management minimizing costs for airspace users: Enabling trajectory options and flexible pre-tactical delay management. *Transportation Research Part B: Methodological*, 134:229–255, 2020. doi: 10.1016/j.trb.2020.02.012.

X. Yang and A. K. Strauss. An approximate dynamic programming approach to attended home delivery management. *European Journal of Operational Research*, 263(3):935–945, 2017. ISSN 03772217. doi: 10.1016/j.ejor.2017.06.034.

X. Yang, A. K. Strauss, C. S. Currie, and R. Eglese. Choice-based demand management and vehicle routing in e-fulfillment. *Transportation Science*, 50(2):473–488, 2016. ISSN 15265447. doi: 10.1287/trsc.2014.0549.

J. Zillies, A. Kuenz, A. R. Schmitt, G. Schwoch, V. Mollwitz, and C. Edinger. Wind optimized routing: An opportunity to improve European flight efficiency? *2014 Integrated Communications, Navigation and Surveillance Conference (ICNS) Conference Proceedings*, pages X3–1–9, 2014.