Contents lists available at ScienceDirect



Transportation Research Part B



journal homepage: www.elsevier.com/locate/trb

The value of flexible flight-to-route assignments in pre-tactical air traffic management

Check for updates

Jan-Rasmus Künnen, Arne K. Strauss*

Chair of Demand Management and Sustainable Transport, WHU - Otto Beisheim School of Management, 56179 Vallendar, Germany

ARTICLE INFO

Keywords: Dynamic pricing Trajectory optimization Demand–capacity balancing

ABSTRACT

In European air traffic management, there are discussions regarding the future role of the network manager (NM): in particular, should the NM be able to assign flights to specific trajectories, should airspace users be allowed to freely choose their preferred trajectory, or something in between? In this paper, we develop a modeling framework that can be adapted to these settings to assess 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 'trajectory 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. Charges are set so as to reward greater flexibility of airspace users with lower charges. The airspace user chooses one of the offered trajectory products according to a choice model that reflects their preferences given, among others, the 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 overall displacement cost. Methodologically, the problem deviates from typical dynamic pricing problems in various major ways, e.g., featuring a boundary condition that we show to be NP-hard as well as fairness and revenue neutrality constraints. The problem is cast in the form of a dynamic program. We exploit a certain structure in the boundary condition to formulate an efficient heuristic. Based on a numerical case study, we find that the use of dynamically priced trajectory products achieves a cost performance close to the one obtained if the NM has a mandate to simply assign flights to trajectories. Therefore, this seems an attractive design for the role of the NM, giving airspace users some choice whilst achieving low overall costs.

1. Introduction

The European air traffic management (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 (creating potential flight delays), 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

* Corresponding author. E-mail addresses: jan-rasmus.kuennen@whu.edu (J.-R. Künnen), arne.strauss@whu.edu (A.K. Strauss).

https://doi.org/10.1016/j.trb.2022.04.004

Received 27 May 2021; Received in revised form 3 March 2022; Accepted 15 April 2022

Available online 4 May 2022

^{0191-2615/© 2022} The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/).

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 level, the network manager (NM) can only implement demand management measures such as assigning delays or rerouting flights 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 CO_2 emissions" (WPG, 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" (WPG, 2019).

A natural question arising in this context is what should be the future role of the NM. In this paper, we focus specifically on developing a modeling framework that can be used to assess the implications of different roles of the NM in pre-tactical ATM. We always assume that the NM is able to decide which airspace runs in which configuration subject to an exogenously given capacity budget; our focus is on demand-side interactions. From a cost minimization perspective, an extreme setting would see the NM empowered to assign trajectories (including delays or reroutings) 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 overall.

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 a 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 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 honoring 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 cost. Displacement cost is the cost incurred by delaying and rerouting flights in order to avoid airspace congestion. 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 conditions, 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 pricing decisions in real-time. Our numerical study provides insights to the debate on the future role of the NM, specifically that dynamically priced trajectory products can achieve an excellent trade-off between AU choice and cost minimization.

The paper is organized as follows: In Section 2 we review current literature on pre-tactical ATM and related fields. In Section 3, we provide a formal description of the DTPP, and Section 4 presents an efficient method to solve the problem for realistic instances. We evaluate dynamic pricing policies based on the proposed method in Section 5 and close with recommendations in Section 6.

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) analyze 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 options 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 subsidize 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 centralized 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 adjust 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 exogenous input to our model. Starita et al. (2020) provide such a capacity planning model; we use a similar formulation for the cost estimation, but solve it heuristically in a different way since we need to make faster dynamic decisions.

Outside the context of ATM, the dynamic pricing of trajectory products in ATM shares some characteristics with the pricing of ride-sharing that has picked up recently due to the success of mobility-on-demand providers such as Uber and Lyft. Our problem is similar in that prices need to be determined in real-time and computing dynamic prices requires solving a hard routing problem in the boundary condition. The main difference is that we need to consider additional pricing constraints that are specific to the ATM environment, such as fairness and revenue neutrality, while the ride-sharing problem features a more traditional revenue-maximization scheme. A good introduction to the ride-sharing problem is given in Hosni et al. (2014). To solve the problem, Chen et al. (2019) develop a Markov decision process formulation and an efficient algorithm that determines revenue-maximizing prices. In a recent work, Alisoltani et al. (2021) have also investigated the potential of ride-sharing to reduce traffic congestion, which is similar in vein to our analysis of whether a stronger role for the NM can reduce displacements of flights in the air traffic network. Furthermore, 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) and 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. The term insertion cost refers to the process of estimating opportunity cost by "inserting" a product (which is to be priced) into various future demand scenarios.

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.

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 booking horizon from t = 1, ..., T, where T is the cut-off time after which no bookings are permitted, and any period t represents a discrete time interval. The length of each interval may vary for each t but is chosen sufficiently small so that at most one request arrives per time period. Within the booking horizon, any AU can submit a flight plan for a flight $f \in F$, 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 min). 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 Section 3.1.3).

We introduce new trajectory product types that the AU can book, which determine how flexibly the NM can decide to route the flight (see Section 3.1.1). 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 3.1.2 describes in detail the modeling of the DTPP as a Markov decision process.

3.1.1. 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 $U = \{1, ..., U^{max}\}$. 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, and R^f represents all possible routes for flight *f*. (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 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 $Z = \{1, ..., n\}$. 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 $R_1^f \subseteq \cdots \subseteq R_n^f = R^f$ 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*).

3.1.2. 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 $X_t \in \chi \subset \mathbb{N}^{|\mathcal{F}| \times n+1}$, which contains in rows $1, \ldots, |\mathcal{F}|$ the product type that has been purchased until *t* (if any) by flight $f \in \mathcal{F}$, and the prices $p^f = (p_z^f)_{z \in \mathbb{Z}}$ that were offered for each product type (of which there are *n*). In particular, the *k*th row $X_{t,k}$ for $k = 1, \ldots, |\mathcal{F}|$ is defined as:

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

Note that $|\mathcal{F}| = n_G + n_N$, where n_G and n_N are the number of all scheduled and non-scheduled flights, respectively. Even if we do not know in advance the non-scheduled flights that will arrive to the booking process, parameter n_N is known since it represents all possible combinations of origin, destination and departure time. It is easy to see that the size of state space *X* grows quickly with the number of product types *n*. However, this does not cause problems since our solution approach in Section 4 does not require iterating through all states *X*; instead we only require the state space to describe the DTPP.

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.,

$$rev_z^f = p_z^f r\bar{e}v^f, \qquad f \in \mathcal{F}, z \in Z, t = 1, \dots, T.$$
(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 $Pr = \{p_i : i = 1, ..., I\}$. 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 cardinality 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.,

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

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 (see Section 3.2.1).

3.1.3 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:

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

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), since we know the average time between arrivals based on average expected flights, and arrivals are independent of one another (i.e., one arrival does not affect the probability of the next). The arrival rate λ^N 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 the arrival of any flight $f \in \overline{F}_t^G$; 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 is updated dynamically after every arrival. In summary, we can estimate the arrival probability of any scheduled or non-scheduled flight as:

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

3.2 Dynamic programming formulation

3.2.1 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:

$$V_{t}(X) = \sum_{f \in \bar{F}_{t}} \lambda_{t}^{f}(\bar{F}_{t}) \min_{p^{f}} \left\{ \sum_{z \in Z} P_{z}(p^{f}) V_{t+1}(X \cup (f, z, p^{f})) \right\} + \left[1 - \sum_{f \in \bar{F}_{t}} \lambda_{t}^{f}(\bar{F}_{t}) \right] V_{t+1}(X)$$

$$= \sum_{f \in \bar{F}_{t}} \lambda_{t}^{f}(\bar{F}_{t}) \min_{p^{f}} \left\{ \sum_{z \in Z} P_{z}(p^{f}) \left[V_{t+1}(X \cup (f, z, p^{f})) - V_{t+1}(X) \right] \right\} + V_{t+1}(X), \quad X \in \chi.$$
(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 $\cot V_{T+1}(X)$. 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 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 remaining uncertainty in the problem setting. Since traffic uncertainty is already described by the materialization of X_{T+1} , W models the uncertainty in capacity provision (due to employee absence or adverse weather). In particular, W represents the actual sector capacities at T + 1 which are modeled based on historic rates (see Section 4.1).

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:

$$V_{T+1}(X) = \mathbb{E}_{X,W}[D(X,W,H)] + \theta^{RN} \epsilon^{RN}(X) + \theta^{FR} \epsilon^{FR}(X), \quad X \in \chi,$$
(3a)

where
$$\epsilon^{RN}(X) := \sum_{f \in F_T} |1 - \hat{p}^f|,$$
 (3b)

$$\epsilon^{FR}(X) := \sum_{f \in F_T} Var(p^f).$$
(3c)

Definition (3b) sets e^{RN} as the absolute deviation between the average collected price and cost-neutral price of 1. 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 ensure 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 (3c) defines e^{FR} as the variance among prices offered for 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. Components e^{RN} and e^{FR} are then used with penalties θ in boundary condition (3a) to model these soft considerations. The penalty terms should be set according to the application at hand, and reflect the choice model $P_z(p^f)$. Parameter θ^{FR} should be set a priori and chosen such that the minimum price variance required to impose any product on the AU (based on the choice model) induces a penalty larger than the maximum benefit from such imposition. In contrast, a breach of *revenue neutrality* is hard to quantify so that θ^{RN} is best defined after initial simulations. As long as a reasonable number of relative prices exist (in Pr), ensuring *revenue neutrality* via (3b) will not impact the objective function value. Even if definitions (3b) and (3c) are constructively similar, they cannot be reduced to one condition because (3b) controls the average price while (3c) controls its variation.

3.2.2 Joint sector-opening and routing optimization

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*, sector capacities *W* and capacity budget *H*, such that total displacement cost is 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 h) 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 lwithin one time interval, based on possibly reduced capacities W. The integrated routing and sector opening problem (IRSOP) is given as:

$$\mathbf{IRSOP}: D(X_t, W, H) = \min_{x, y} \sum_{f \in F_t} \sum_{r \in R_\tau^f} d_r^f x_r^f$$
(4a)

s.t.
$$\sum_{u} \sum_{c \in C^{a}} \bar{h}_{ac} y_{acu} \le H_{a} \qquad a \in A$$
(4b)

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

$$(4c)$$

$$\sum_{c \in C^a} y_{acu} = 1 \qquad a \in A, u \in U$$
(4d)

$$\sum_{e \in R_z^f} x_r^f = 1 \qquad \qquad f \in F \tag{4e}$$

$$f \in F, r \in R_z^f \tag{4f}$$

$$y_{acu} \in \{0, 1\} \qquad a \in A, c \in C^a, u \in U.$$
(4g)

The objective function minimizes total flight displacement cost over all flights and routes. Constraint (4b) defines the feasible configurations based on the capacity budget H_a for each airspace. Constraint (4c) 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 ; otherwise the left-hand side reduces to 0 and the constraint holds. Constraint (4d) ensures that each airspace operates under one configuration at any time, (4e) ensures that exactly one route is assigned to each flight, and (4f)–(4g) model the binary condition for our decision variables.

 $x_r^f \in \{0, 1\}$

3.2.3 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) that this online decision can be made as follows:

$$p^{f}(X) = \arg\min_{p^{f}} \sum_{z \in \mathbb{Z}} P_{z}(p^{f}) \Delta_{(j,p)} V_{t+1}(X).$$
(5)

We set prices p_i^f for flight f such that expected opportunity cost $\Delta_{(j,p)}V_{i+1}(X)$ is minimized. The pricing policy is dynamic in that we adjust the price vector p^f based on incoming bookings over time. However, note that the real-time control policy in (5) does not depend on time t. That is, we do not explicitly include a timing incentive that would, for instance, reward earlier bookings.

Determining optimal prices in (5), 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 which are defined by our action space with cardinality I^n . In our experiments, the number of price points I and product types n is small enough so that we can in fact fully enumerate the price vectors and their respective objective value in (5). If full enumeration is not possible in real-time, one of the assortment optimization techniques discussed in Strauss et al. (2018) can be applied. By exploiting the structure of the underlying choice model, these techniques allow finding optimal solutions in polynomial time. For instance, Dong et al. (2009) prove the concavity of choice probabilities under the multinomial logit (MNL) model and derive a quick analytical solution.

In summary, the key to making optimal decisions is to quantify the opportunity cost, which in turn depends on the value function. Since determining the value function exactly 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 estimate opportunity cost offline so that they can be used in the real-time decision policy.

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 $S \in S(X)$. In particular, a scenario describes a forecast of flight arrivals and actual capacity for departure day. Here, $S(X_i)$ denotes the population of scenarios for $X_i \in \chi$; it depends on X_i 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 $t \le T$ under scenario $S \in S(X_i)$. 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:

$$V_{t}(X_{t}) \approx \hat{V}_{t}(X_{t}) := \sum_{S \in S(X_{t})} \min_{\pi^{S}} \frac{V_{T+1}(X_{T+1}^{\pi^{S},S})}{|S(X_{t})|}.$$
(6)

That means, we approximate the value function at *t* with the expected value at T + 1, evaluated over $S \in S(X_t)$, 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:

$$\lim_{|S| \to \infty} \sum_{S \in S(X_t)} \min_{\pi^S} \frac{V_{T+1}(X_{T+1}^{\pi^3, S})}{|S(X_t)|} \le V_t(X_t) \le \lim_{|S| \to \infty} \sum_{S \in S(X_t)} \frac{V_{T+1}(X_{T+1}^{\pi, S})}{|S(X_t)|}.$$

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:

$$V_{t}(X) = \min_{\pi} \mathbb{E}_{S}[V_{T+1}(X_{T+1}^{\pi,S})] = \min_{SAA} \min_{\pi} \lim_{|S| \to \infty} \sum_{S \in S(X_{t})} \frac{V_{T+1}(X_{T+1}^{\pi,S})}{|S(X_{t})|} \ge \lim_{|S| \to \infty} \sum_{S \in S(X_{t})} \min_{\pi^{S}} \frac{V_{T+1}(X_{T+1}^{\pi^{3},S})}{|S(X_{t})|}.$$

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:

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

Using the approximation in (6), we obtain for the opportunity cost:

$$\Delta_{(j,p)}V_{t+1}(X_t) \approx \Delta_{(j,p)}\hat{V}_{t+1}(X_t) = \sum_{S \in S(X_t)} \frac{\Delta_j D(X_{T+1}^S, W^S, H)}{|S(X_t)|} + \theta^{RN} \Delta_p \epsilon^{RN} + \theta^{FR} \Delta_p \epsilon^{FR}.$$
(7)

To obtain (7) we first replace $V_{T+1}(X_{T+1}^{\pi^S,S})$ in (6) with the respective displacement and penalty cost considerations in V_{T+1} and then determine opportunity cost by computing the change in displacement and penalty cost with regards to product j and price p. The change in displacement cost $\Delta_j D(X_{T+1}^S, W^S, H)$ is computed as the cost of "inserting" product j = (f, z) to final bookings X_{T+1}^S (under scenario S), which is why we refer to it as the insertion cost of product j. The changes in penalty cost $\theta^{RN} \Delta_p \epsilon^{RN}$ and $\theta^{FR} \Delta_p \epsilon^{FR}$ only depend on price p (since information on product j or scenario S do not influence the revenue neutrality or fairness of state X_{T+1}), so we do not need to evaluate these costs over $S \in S(X)$. The quality of the approximation in (7) 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 Section 4.1;
- 2. Determine routing of flights with lowest cost $D(X_{T+1}^S, W^S, H)$ for each sample, see Section 4.2;
- 3. Determine $\cot \Delta_j D(X_{T+1}^S, W^S, H)$ as well as penalty parameters $\Delta_p e^{RN}$ and $\Delta_p e^{FR}$ for adding product *j* at price set p^f to the routing of any sample (X_{T+1}^S, W^S) , see Section 4.3.

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}_t^G). 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 $S = (X_{T+1}^S, W^S)$ that differ with regard to the two inputs mentioned above. For X_{T+1}^S , 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 $\mu_F - |F_t| - |\bar{F}_t^G|$, 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 $r \in \mathbb{R}^f$, 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.

1

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 sector capacities W^S can be modeled in various ways as long as the resulting capacity levels reflect historic fluctuations in the airspace. In our case, we choose at random one elementary sector in each airspace, and reduce the capacity of this sector (and all collapsed sectors containing it) in line with historic rates. In practice, these rates can be determined based on past ATM regulations which are imposed on airspaces and reduce its nominal sector capacities; they should be observed over a period of over one year to allow for a fair representation of extreme events (such as strikes or abnormal weather conditions).

4.2 Metaheuristic approach for routing problem

In a first step, we need to determine the routing with lowest displacement cost $D(X_{T+1}^S, W^S, H)$ in (4a) for each sample scenario (X_{T+1}^S, W^S) , which requires solving the IRSOP. However, this problem is \mathcal{NP} -hard, as shown below.

Theorem 1. The IRSOP described by (4a)-(4f) 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 (4b) 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 (4b),(4d)–(4g) become redundant, and the remaining problem changes to:

$$D(X_t, W, H) = \min_{x} \sum_{f \in F_t} \sum_{r \in R_z^f} d_r^f x_r^f$$

s.t.
$$\sum_{f \in F_t} \sum_{r \in R_z^f} \sum_{e \in E^l} b_{freu} x_r^f \le \mathcal{K}_l$$

$$C' \in C', l \in L^{c'}, u \in U$$

$$\sum_{r \in R_z^f} x_r^f = 1$$

$$x_r^f \in \{0, 1\}$$

$$f \in F, r \in R_z^f.$$

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). \Box

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.

4.2.1 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 $k_{acu} := \sum_{l \in L^c} \left(\sum_{e \in E^l} \sum_{f \in F_l} \sum_{r \in R_{z'}^f} b_{freu} x_r^f - \mathcal{K}_l \right)^+$, where $x^+ := \max\{x, 0\}$. Configuration set C' can then be determined with the following configuration ILP:

 $\mathbf{CILP}: \min_{y} \sum_{a,c,u} k_{acu} y_{acu}$ (8a)

s.t.
$$\sum_{u} \sum_{c \in C^{a}} \bar{h}_{ac} y_{acu} \le H_{a}$$

$$\sum_{c \in C^{a}} y_{acu} = 1$$

$$a \in A, u \in U$$
(8b)
(8b)
(8c)

$$a \in A, c \in C^a, u \in U.$$
(8d)

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).

4.2.2 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 $\bar{k}_l = \sum_{f \in F_i} w_{fr'l}$ 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,

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

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: $r'_f := \operatorname{argmin}_{r \in \mathbb{R}^f} d^f_r$ 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 $l^* := \operatorname{argmax}_l \bar{k}_l$
- 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 $f \in F_t$, $r \in \mathbb{R}^f$.
- 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:

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

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., lowcost) 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.

4.3 Insertion heuristic to determine opportunity cost

Now that we have developed suitable routings for sample scenarios (X_{T+1}^S, W^S) , we want to compute opportunity cost for selling any product j = (f, z) at price set p^f , where $f \in \overline{F}_t$ 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 $\Delta_j D(X_{T+1}^S, W^S, H)$ of adding product j to the final routings for sample (X_{T+1}^S, W^S) . We have

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

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_{T+1}^S, W^S) , we then determine route $r' \in R_z^f$ that generates the lowest increase in displacement cost while keeping the resulting solution feasible, i.e., the route with lowest $\delta_r^{f,S}$ where

$$\delta_r^{f,S} = \begin{cases} d_r^f & \text{if } f \in F^G \text{ and } \bar{k}_l \leq 1, l \in L' \\ d_r^f & \text{if } f \in F_t \setminus F_t^G \text{ and } \bar{k}_l + w_{frl} \leq 1, l \in L' \\ M & \text{otherwise.} \end{cases}$$

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 $S \in S(X)$) 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 $\Delta_j D(X_{T+1}^S, H) = \delta_{r'}^{S,f}$ is then used in (7) 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 Section 3.2.1. To ensure *revenue neutrality* and *fairness*, we compute:

$$\Delta_{p} \epsilon^{RN} = |1 - \hat{p}^{f}|$$

$$(9)$$

$$A \epsilon^{FR} = Var(n^{f})$$

$$(10)$$

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 Section 5.1 we describe the settings and corresponding decision policies which we evaluate, in Section 5.2 we discuss the simulation study and evaluation process, and in Sections 5.3-5.4 we report our results.

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 Section 3.1.1) 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, which we introduce in Section 5.2.1) 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 (5), where we compute expected opportunity cost as:

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



Fig. 1. Binary logit function to model AU choice.

• Foresight dynamic pricing (FD): In contrast to HD, insertion cost is determined based on the foresight approach described in Section 4.3, where each scenario reflects a full schedule for cut-off time T + 1. We compute opportunity cost as follows:

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

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 Section 4. Note that all prices are quoted as relative prices with regards to the benchmark price on a certain origin–destination pair (see Section 3.2.1), so they automatically reflect the route characteristics of a flight.

5.2 Simulation study

5.2.1 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_t^f based on the respective pricing policy. For all dynamic pricing policies, we compute opportunity cost from a fixed number of scenarios $S \in S(X)$. Each scenario S consists of a set of flights and product types for the end of the booking horizon as well as sector capacities W, and is created using the sampling strategy in Section 4.1. Recall that the offline process of determining routings for all $S \in S(X)$ is the most time-consuming part in computing opportunity cost. Therefore, we need to decide carefully on two parameters that impact the solving time, and solution quality, of this process: the frequency with which the offline process should be re-run (to reflect the latest bookings X_t in S(X)), and the number of scenarios $S \in S(X)$ to use in (7) to compute opportunity cost. We decide to re-run the offline process after every 10 flights that have arrived to the booking process, and to use 20 scenarios for each computation. Any higher frequency of updates or higher number of scenarios did not notably improve the quality of estimates in initial simulation runs. The scenarios are randomly selected from S(X) and are fixed across policies, but vary for each simulation run.

In our simulation study we choose to offer two product types, one flexible and one direct product: $Z = \{flex, dir\}$. Furthermore, we allow the prices p_z^f to be chosen among $Pr = \{0.9, 0.91, \dots, 1.4\}$, which contains a total of 51 scaling factors (i.e., I = 51). 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, $v = p_{flex}/p_{dir}$. 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 v (which we set to 0.85) (see Fig. 1):

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

Given all product choices for t = 1, ..., 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 2000 runs for each policy.

5.2.2 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 Fig. 2. 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) 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



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

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 $S \in S(X)$, 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 min). 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 $R_{dir}^{f} := r_{f}^{*}$ (where r_{e}^{*} is the shortest route of flight f) and $R_{flor}^{f} := R^{f}$.

We consider an operating time interval of 2 h, with configuration changes allowed every 30 min. 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 is derived from historic ATM regulations in 2016, in which weather events occurred on around 5% of days across 15 ACCs in Europe and impacted 30% of capacity, and employee absence occurred for another 5% and impacted 10% of capacity, on average. The reductions apply across the 2 h operating time interval. Finally, 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). We provide the full case study dataset in Tables 3–6 to allow results to be replicated.

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 $D(X_{T+1}^S, W^S, H)$ on a total of 20 scenarios, where X_{T+1}^S and W^S are generated as described in Section 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 on 20 instances, and are summarized in Table 1 (the full results on all instances are reported in Tables 3–6). 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.

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. Fig. 3 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 1 determined through either the ILP or heuristic (see (7)). 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 median absolute error of 14.5, which is

Table 1

Summary of simulation results of routing model (n = 20 instances).

	ILP			Heur.				
	Cost	Time (s)	N. a.	Cost	Time (s)	N. a.	Cost gap	
ø	26,023	3022	4.4%	28,852	2.2	5.6%	11.3%	

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



Fig. 3. Opportunity cost of ILP vs. heuristic on 200 flights.

slightly more than 1% of average opportunity cost of 1250. The median absolute error is chosen over the mean absolute error (MAE) as measure of performance since a few large errors inflate the MAE, as can be seen in Fig. 3.

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 s, compared to 3022 s \approx 50 min for the exact approach. To understand how the solution time scales for larger instances, recall that the heuristic has complexity $O(m(n - g)^2 + mn)$ 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. 1350) and *g* the number of flights (approx. 150). If we assume that the solving time increases linearly with the complexity of the heuristic (which is realistic if the algorithm has reached asymptotic behavior for the analyzed case study), a realistic-sized instance in European airspace (25,000 daily flights, 7 route options each, and 2000 sectors opened) would require around 20 days to solve with the heuristic. It is important to note that each scenario can always be evaluated independently (since the solution of one scenario does not influence the others) so that the process of computing opportunity cost can be parallelized across scenarios.

In its current form, the heuristic is still too complex to solve a realistic-sized instance sufficiently fast. However, there are two ways to reduce the solution time further (e.g., to allow for daily updates of opportunity cost), which warrant further research: First, the number of flights may be limited to those frequently affected by reroutings. 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. If these flights can be identified *a priori*, they can be excluded from the optimization by assigning them to their shortest route. Second, the routing problem may be decomposed. For instance, the day of operation may be split into distinct time intervals (e.g., morning, afternoon and evening), and the problem solved for each interval separately. This decomposition would substantially reduce solution time (due to the impact on both *n* and *m*), but it would also impact the problem structure and the solution quality.

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. The following Table 2 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. We want to emphasize at this point that the experiments are conducted to illustrate the value of the methodological approach to determine dynamic prices, and that the results are specific to the experimental setup. The limitations to generalizing these results are discussed in Section 5.4.3.

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 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 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. The FD also significantly outperforms the static pricing policy (FS), showing that

J.-R. Künnen and A.K. Strauss

Table 2

Simulation results of pricing policies over 2000 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.



Fig. 4. Displacement cost distribution across all runs, n = 2000 (runs).

there is further value in dynamically adjusting prices for flexible products. Interestingly, the FS policy reports lower cost than the HD policy, which 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 assigned 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 both dynamic pricing policies (HD and FD) cover quite exactly 100.0% of capacity cost, which suggests that requirement (9) is 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 a price of 1 for any 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.

5.4.1 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 s, and by the FD policy in 32 s. 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 s for FD does not pose any problems for implementation, it increases roughly quadratically with the size of the case study (Section 5.3). As long as the offline calculation, parallelized across scenarios, can be carried out in 24 h or less, the FD policy is preferred due to its superior performance. If computing time renders a daily update of the FD policy infeasible, static pricing (FS) is preferred over the HD and FD policies. The online calculation (i.e., determining optimal price sets) is performed in 0.07 and 0.2 s for the HD and FD policies, respectively, which renders both policies suitable for real-time application. Most importantly however, the solution time of the online process does not depend on the size of the case study so it remains sufficiently fast for larger networks.

To analyze how effectively each policy reduces displacement cost, we compare in Fig. 4 the distribution of total displacement cost for the NMd, FD and AUd policies across 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.

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 1475 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.



Fig. 5. Distribution of average displacement cost by policy, n = 200 (flights).



Fig. 6. Distribution of average displacement cost on FD policy by aircraft type, n = 200 (56 small, 108 medium, 36 large).

5.4.2 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 analyze how displacement cost are distributed among flights in Figs. 5 and 6. In contrast to Fig. 4, where we compare total displacement cost between runs, Figs. 5–6 compare average displacement cost per flight between all 200 flights in the case study. Similar to the earlier observation, Fig. 5 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 1200 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.

Fig. 6 compares the distribution of displacement cost under the FD policy for different aircraft types. We see that 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.

Finally, Fig. 7 shows how the average prices offered under the FD policy differ across time and for different aircraft types. We analyze only scheduled flights since they form part of every traffic scenario and as such are more likely to feature several observations for a single booking time *t*. We find that relative prices are highest for small aircraft, followed by medium and large aircraft, and that this observations is specifically pronounced for the *direct* product. This is because small aircraft are more likely to be rerouted due to their lower displacement cost (see Fig. 6) which makes the *direct* product less desirable and thus more expensive for these flights. We also see that the price fluctuations are highest for large aircraft, which is likely due to the fact that the large displacement cost of these flights include larger price differences (depending on whether the flights crosses congested airspace or not). In response, AUs operating large aircraft may benefit from checking prices more frequently, particularly if they intend to purchase the *direct* product.

Based on the prices offered for both product types, we can also test whether the *fairness condition* (see Section 3.2.1) prevents the NM from choosing prices such that one product type is effectively "imposed" on the AU. In fact, the average price offered for the *flexible* 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.



Fig. 7. Average price offered under FD policy by aircraft and product type, n = 120 (scheduled flights).

 Table 3

 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%	
ø	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.

5.4.3 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.

6 Conclusions

We demonstrate the value of dynamic pricing to steer demand in pre-tactical ATM. European ATM suffers from severe demandcapacity 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 particularly

Table 4

Overview of routes for case study.

OD	(Elementa	ary sector, time	period)				Delay	Displacen	nent 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)	(10.0)	(6 1)	(7.1)	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	12/	162	1
2	(2, 1)	(3, 1)	(8, 2)	(10, 2)	(6, 1)	(7, 1)	1	48	90	212	1
2	(2, 1)	(3, 1)	(8, 2)	(10, 2)	(0, 1)	(7, 1)	2	120	250	515	1
2	(2, 1)	(3, 1)	(0, 2)	(10, 2)	(0, 1)	(7, 1)	3	204	430	888	1
2	(2, 1)	(3, 1)	(8, 2)	(10, 2)	(0, 1)	(7, 1)	5	453	1004	1275	1
2	(2, 1)	(3, 1)	(8, 2)	(10, 2)	(6, 1)	(7, 1)	5	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 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	12/5	1
5	(4, 2)	(0, 4)	(1, 2)	(0, 0)	(0, 1)	(0, 1)	6	611	1390	1740	1
0	(7, 1)	(0, 1)	(10, 2)	(8, 2)	(3, 1)	(2, 1)	0	150	0	0	0
6	(7, 1)	(0, 1)	(1, 2)	(0, 2)	(3, 1)	(2, 1)	0	152	∠80 107	355	1
0	(7, 1)	(0, 1)	(1, 2)	(0, 2)	(3, 1)	(2, 1)	0	09 47	12/	102	1
6	(7, 1)	(0, 1)	(10, 2)	(0, 2) (8 2)	(3, 1)	(2, 1)	1 2	47 120	90 926	313	1
6	(7, 1)	(0, 1)	(10, 2)	(8, 2)	(3, 1)	(2, 1)	2	204	250	560	1
0	(7, 1)	(0, 1)	(10, 2)	(0, 2)	(3, 1)	(2, 1)	J	204	430	500	T

(continued on next page)

Table 4 (continued).

OD	(Elementa	ary sector, time	e period)				Delay	Displacen	ient cost		z
	1	2	3	4	5	6		Small	med.	Large	
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 5

Overview	of	flights	for	case	study.
----------	----	---------	-----	------	--------

Flight ID	OD	Aircraft type	Dep. time	Flight type
FO	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.

challenging since it requires solving a hard routing problem in the terminal condition, which itself is subject to uncertainty around arriving flights 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 in the case

Table 6				
O	~£	a a m fi au una ti a m a	6	 ater day

Sector ID	Airspace	Conf. ID	Elementary	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	Т	C5	4	5		18
P7	Т	C6	4			18
P8	Т	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.

study. 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. Future research on the topic may consider, next to an application of the methodology to a larger network, the introduction of a timing incentive that would reward earlier bookings with a lower price. This would give the network manager more time to suggest alternative routes before departure.

CRediT authorship contribution statement

Jan-Rasmus Künnen: Writing – original draft, Software, Conceptualization, Methodology, Formal analysis, Investigation. Arne K. Strauss: Supervision, Conceptualization, Methodology, Writing – review & editing, Funding acquisition.

Acknowledgments

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 program 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.

References

- Alisoltani, N., Leclercq, L., Zargayouna, M., 2021. Can dynamic ride-sharing reduce traffic congestion? Transp. Res. B (ISSN: 0191-2615) 145, 212–246. http://dx.doi.org/10.1016/j.trb.2021.01.004.
- Barnhart, C., Fearing, D., Odoni, A., Vaze, V., 2012. Demand and capacity management in air transportation. EURO J. Transp. Logist. (ISSN: 21924376) 1 (1–2), 135–155. http://dx.doi.org/10.1007/s13676-012-0006-9.

Bolić, T., Castelli, L., Rigonat, D., 2017. Peak-load pricing for the European air traffic management system using modulation of en-route charges. Eur. J. Transp. Infrastruct. Res. (ISSN: 15677141) 17 (1), 136–152. http://dx.doi.org/10.18757/ejtir.2017.17.1.3184.

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

Castelli, L., Labbé, M., Violin, A., 2013. A network pricing formulation for the revenue maximization of European air navigation service providers. Transp. Res. C (ISSN: 0968090X) 33 (June 2018), 214–226. http://dx.doi.org/10.1016/j.trc.2012.04.013.

Chen, M., Shen, W., Tang, P., Zuo, S., 2019. Dispatching through pricing: Modeling ride-sharing and designing dynamic prices. In: Proceedings of the 28th International Joint Conference on Artificial Intelligence. pp. 165–171.

Dong, L., Kouvelis, P., Tian, Z., 2009. Dynamic pricing and inventory control of substitute products. Manuf. Serv. Oper. Manage. 11 (2), 317-339.

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

Hosni, H., Naoum-Sawaya, J., Artail, H., 2014. The shared-taxi problem: Formulation and solution methods. Transp. Res. B (ISSN: 0191-2615) 70, 303–318. http://dx.doi.org/10.1016/j.trb.2014.09.011.

Ivanov, N., Jovanović, R., Fichert, F., Strauss, A., Starita, S., Babić, O., Pavlović, G., 2019. Coordinated capacity and demand management in a redesigned air traffic management value-chain. J. Air Transp. Manage. (ISSN: 09696997) 75 (July 2018), 139–152. http://dx.doi.org/10.1016/j.jairtraman.2018.12.007.

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

- Jovanović, R., Tošić, V., Čangalović, M., Stanojević, M., 2014. Anticipatory modulation of air navigation charges to balance the use of airspace network capacities. Transp. Res. A (ISSN: 09658564) 61, 84–99. http://dx.doi.org/10.1016/j.tra.2014.01.005.
- Martello, S., 1990. Knapsack Problems: Algorithms and Computer Implementations. In: Wiley-Interscience Series in Discrete Mathematics and Optimization, John Wiley & Sons Ltd..
- Melgosa, M., Prats, X., Xu, Y., Delgado, L., 2019. 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, http://dx.doi.org/10.1109/DASC43569.2019.9081803.
- Moser, M., Jokanovic, D.P., Shiratori, N., 1997. An algorithm for the multidimensional multiple-choice knapsack problem. IEICE Trans. Fundam. Electron. Commun. Comput. Sci. 80 (3), 582–589.
- Mukherjee, A., Hansen, M., 2009. A dynamic rerouting model for air traffic flow management. Transp. Res. B 43 (1), 159–171. http://dx.doi.org/10.1016/j.trb. 2008.05.011.
- Parlar, M., Rodrigues, B., Sharafali, M., 2018. Event-based allocation of airline check-in counters: a simple dynamic optimization method supported by empirical data. Int. Trans. Oper. Res. (ISSN: 14753995) 25 (5), 1553–1582. http://dx.doi.org/10.1111/itor.12332.
- Pisinger, D., 1995. A minimal algorithm for the multiple-choice knapsack problem. European J. Oper. Res. (ISSN: 03772217) 83 (2), 394–410. http://dx.doi.org/10.1016/0377-2217(95)00015-I.
- Rambha, T., Boyles, S.D., 2016. Dynamic pricing in discrete time stochastic day-to-day route choice models. Transp. Res. B 92, 104–118. http://dx.doi.org/10. 1016/j.trb.2016.01.008.
- Ranieri, A., Castelli, L., 2008. Pricing schemes based on air navigation service charges to reduce en-route ATFM delays. In: Third International Conference on Research in Air Transportation.
- Starita, S., Strauss, A.K., Fei, X., Jovanović, R., Ivanov, N., Pavlović, G., Fichert, F., 2020. Air traffic control capacity planning under demand and capacity provision uncertainty. Transp. Sci. 54 (4), 882–896. http://dx.doi.org/10.1287/trsc.2019.0962.
- Starita, S., Strauss, A.K., Jovanovic, R., Ivanov, N., Fichert, F., 2016. Maximizing ATM cost-efficiency by flexible provision of airspace capacity. In: SESAR Innovation Days.
- Steer Davies Gleave, 2015. Policy Options for the Modulation of Charges in the Single European Sky. Technical Report April.
- Strauss, A.K., Klein, R., Steinhardt, C., 2018. A review of choice-based revenue management: Theory and methods. European J. Oper. Res. 271 (2), 375–387. http://dx.doi.org/10.1016/j.ejor.2018.01.011.

WPG, 2019. Report of the Wise Persons Group on the Future of the Single European Sky 2019. Report, European Commission.

- Xu, Y., Dalmau, R., Melgosa, M., Montlaur, A., Prats, X., 2020. A framework for collaborative air traffic flow management minimizing costs for airspace users: Enabling trajectory options and flexible pre-tactical delay management. Transp. Res. B 134, 229–255. http://dx.doi.org/10.1016/j.trb.2020.02.012.
- Yang, X., Strauss, A.K., 2017. An approximate dynamic programming approach to attended home delivery management. European J. Oper. Res. (ISSN: 03772217) 263 (3), 935–945. http://dx.doi.org/10.1016/j.ejor.2017.06.034.
- Yang, X., Strauss, A.K., Currie, C.S., Eglese, R., 2016. Choice-based demand management and vehicle routing in e-fulfillment. Transp. Sci. (ISSN: 15265447) 50 (2), 473–488. http://dx.doi.org/10.1287/trsc.2014.0549.