J Revenue Pricing Manag
https://doi.org/10.1057/s41272-018-0147-z

PRACTICE ARTICLE

Dynamic pricing of airline offers

Thomas Fiig[1] · Remy Le Guen[2] · Mathilde Gauchet[2]

Received: 22 October 2017 / Accepted: 9 March 2018

Abstract Airlines have started to focus on expanding their product offerings beyond flights to include ancillary products (e.g., baggage, advance seat reservations, meals, flexibility options), as well as third-party content (e.g., parking and insurance). Today, however, offer creation is rudimentary, managed in separate processes, organizations, and IT systems. We believe the current approach is inadequate and that the key to profitability is to manage offers consistently in an integrated Offer Management System (OMS). However, realizing this vision, will require significant advancements in the science of pricing and in distribution. The entire scope of an OMS cannot be covered in a single paper. Hence, to provide depth, we will focus on what we believe is one of the most critical components of the OMS—Dynamic Pricing of airline offers. Finally, we discuss various industry initiatives that will enable deployment of Dynamic Pricing of the flight product alone or the broader scope of OMS.

Keywords Revenue Management System · Offer Management System · Dynamic Pricing · Ancillary Price Optimization · Machine Learning · Distribution · NDC

✉ Thomas Fiig
Thomas.Fiig@Sca.Amadeus.com

1 Amadeus Airline IT, Lufthavnsboulevarden 14, 2770 Kastrup, Denmark

2 Amadeus Airline IT, Amadeus S.A.S., 485, Route du Pin Montard, 06902 Sophia Antipolis Cedex, France

Introduction and motivation

Profitability is tough in the airline industry. An IATA economic briefing (2013) concluded that the return on invested capital for airlines was the lowest among 30 industries in the comparison study. To improve profitability, airlines have started to focus on expanding their product offerings beyond flights to include ancillary products, such as baggage, advance seat reservations (ASR), meal, and flexibility options, as well as third-party content.

However, offer construction is quite rudimentary today. The flight product is carefully priced through the application of revenue management systems (RMS), which have about 40 years of advancement behind them. Yet little or no consideration is commonly given to how ancillary products are selected or, for that matter, priced. Today, this is handled by two distinct processes, RMS and Merchandising, which are managed under separate organizations through different IT systems.

RMS is responsible for optimizing revenue from flight products alone, while Merchandising is responsible for expanding the shopping basket through upsell, cross-sell, and through selling ancillary products. Also, limited attention is paid to providing personalized and more relevant offers.

We believe the current approach is inadequate and that the key to profitability is to manage offers consistently in an integrated Offer Management System (OMS). This is not possible today, given the limitations of legacy IT systems. Legacy distribution systems delegate airline control of offer creation to content aggregators, such as global distribution systems (GDSs), building itineraries and pricing them on the basis of filed fares and fare rules. This means that airlines are unaware of the customer's identity and have limited control over the offer construction. These

Published online: 12 April 2018

limitations prevent the airline industry from adopting modern e-commerce and enjoying its potential benefits (Popp 2016). At a more detailed level, these limitations are described below.

Description and limitations of existing RMS

Incorrect view and valuation of the flight product

- In a shopping session, availability and pricing decisions are for flights only. The value of ancillary products is disregarded.
- Even considering flights only, a precise valuation would require continuous prices. This is not possible with the 26 Reservation Booking Designator (RBD)-based inventory logic that exists today.
- The net revenue received by the airline from selling a flight product is uncertain, as multiple fare products with different fares are filed in the same RBD.
- Similar uncertainty exists with regard to interline itineraries, the value of which is prorated across interline partners according to special prorate agreements (SPAs). RMS neglects the variety and sophistication of SPAs, relying instead on approximations.

Incorrect view of demand and willingness to pay

- Demand forecasting and pricing work in different dimensions, leading to erroneous forecasting of demand and willingness to pay (WTP):
  - The RMS' demand forecasting utilizes dimensions to support control of the flight resources on the basis of traffic flow, point of sale, and booking class.
  - The airline's pricing department utilizes dimensions to support differentiation of a customer's WTP, based on rules and restrictions such as customer eligibility, return trip restrictions, minimum and maximum stay durations, stopovers, Saturday and Sunday night stay requirements, advance-purchase restrictions, and combinability.

No optimization of the total product offer

- The total set of products sold by the airline is neither optimized nor priced together, but independently.

No personalization or merchandizing

- All customers receive the same price for the same products.
- The sales display cannot be tailored to a customer in a way that affects his or her purchase behavior.

Contribution

In this paper we will formulate our vision for an OMS, and discuss how we may overcome many of the limitations that exist today. However, the entire scope of an OMS cannot be covered in a single paper. To provide depth, we will focus on what we believe is one of the most critical components of an OMS—dynamic pricing (DP).

Additionally, we discuss how the changing distribution landscape, shaped by such initiatives as IATA's New Distribution Capability (NDC) and a global digital transformation, will dramatically influence the airline OMS.

To our knowledge, no prior published work addresses these questions, nor are we aware of any implementation that includes dynamic pricing of airline offers.

Organization of the paper

The paper is structured as follows: in Sect. "The Offer Management System," we define the airline offer, formulate our vision for the OMS, provide an overview of the OMS and its main components, and discuss the scientific challenges of pricing a general airline offer. In Sect. "Dynamic pricing of offers," we discuss DP of airline offers, the underlying assumptions, and application to both base and ancillary products. In Sect. "A distribution landscape in motion," we discuss various industry initiatives that will enable deployment of DP or the broader OMS, the distribution channels and their interactions with the OMS, and the technology required. In Sect. "Future research directions," we offer direction for future research, and finally in Sect. "Conclusion," we provide our conclusions.

The Offer Management System

Description of an offer

We apply the following *definition*, in accordance with IATA (2016), "An *offer* is a proposal by an airline to a customer for a defined *set of products* (flights and/or flight related or non-flight related ancillary products) in response to a shopping request received from a seller, possibly via an aggregator." Shopping requests may be sent with individual traveler data (personalized) or sent without any specific traveler data (anonymous). The shopping response may be customized based on the traveler information that is passed in the request (see also Sect. "Scientific challenges of pricing an offer").

Offers may not be altered, nor can elements of one offer be combined with elements of another offer. The offer will have certain rules and conditions attached to it, for example how long the offer is valid for, by what time payment must be made, how long inventory is guaranteed for, etc.

An offer consists of one or more *offer items*. An offer item consists of two parts: a *set of products*[1] and a *price;* bundled together in unity. The final property of the offer item we would like to mention is its status, *mandatory* or *optional*. During the shopping process an offer can be customized by selecting/deselecting optional offer items. Consequently, the total offer price is computed by summing the selected offer items prices.

To simplify the description of offers in this paper, we will use the following terminology. We denote by *product offer* and *price offer*; the sets of products and the prices of the constituting offer items, respectively. Finally we denote the online processes that determine the product offer and price offer, *dynamic offer construction* and *dynamic pricing*, respectively.

Before we can understand the complexities of how to dynamically price an offer, let us first consider a general product offer.

Product set of the airline offer

Figure 1 illustrates a general airline offer composed of mandatory and optional offer items. The individual product sets that make up the offer items are described below. Offer items are marked by dashed boxes. The final offer may be customized by selecting/deselecting among the optional offer items.

- *Flight products*—fare families with connecting flights (*AL-FLT1, AL-FLT2*): bundles consisting of flight seats, in-flight service, and privileges/ancillary products (free of charge) attached with conditions and restrictions (shown as *business, premium economy,* and *economy* products).
- *Hotel rooms*—fare families with "connecting rooms" (*Hotel A—Day 1*, *Hotel A—Day 2*): airlines' own hotel properties or third-party hotels (shown as *sea-view and street-view*).
- *Interline flight connections* (*OAL-FLT*).
- *Base product*—mandatory offer item that includes the flight product. O*ne and only one* of the flight products (economy, premium economy, and business fare family) must be selected. Similarly for the hotel stay.
- *Bundles of ancillaries*—optional offer item targeted to specific customer segments, such as a baggage pack (shown as *carry-on bag*, *1st bag*, and *2nd bag*) or a flexibility pack (shown as *unlimited changes*, *name change*, or *different itinerary*). Other possibilities could be a comfort pack (not shown), consisting of *priority boarding, middle seat free,* and/or *ASR*.
- *Ancillaries offered a la carte*—optional offer items (e.g., *fare lock, cancelation, fast track,* or *priority boarding*), may be available only on an a la carte basis, while other ancillaries (e.g., *carry-on bag, 1st bag, 2nd bag*) may also be available as bundles.
- *Third-party content*—optional offer items (e.g., *insurance*, *parking*, *hotels*, *rental car*).

Description of an Offer Management System

Vision

We believe that:

- *The OMS should dynamically construct and dynamically price offers, considering both customer and contextual information*.
- *The distribution environment should provide channel consistency with accurate content, built and controlled centrally by the airline*.

Having a detailed view of the customer's preferences and WTP during the shopping session allow airlines to differentiate among customers, thus making their offers more relevant and improving the customer experience. This will enable airlines to improve their revenue performance. Further, channel consistency will eliminate the confusion and mistrust that exist today among consumers about offer transparency.

Components of the OMS

The business logic of OMS relies on five main modules—*Content Management*, *Customer Segmentation*, *Dynamic Offer Construction*, *Dynamic Pricing,* and *Merchandizing*—that are briefly described below. Figure 2 shows how OMS processes an incoming shopping request through each of these modules and responds with an offer (or multiple offers). The DP component, which is highlighted in the figure, is the primary focus of this paper.

Content management The objective of content management is to provide a catalog of the entire set of products (bundled and a la carte) that the airline can sell, as

[1] The IATA (2016) terminology is *set of services*. The terms product and service seems to be used interchangeable in the airline industry. Hence to simplify the description of offers in this paper we will use the term products to collectively describe products and services.

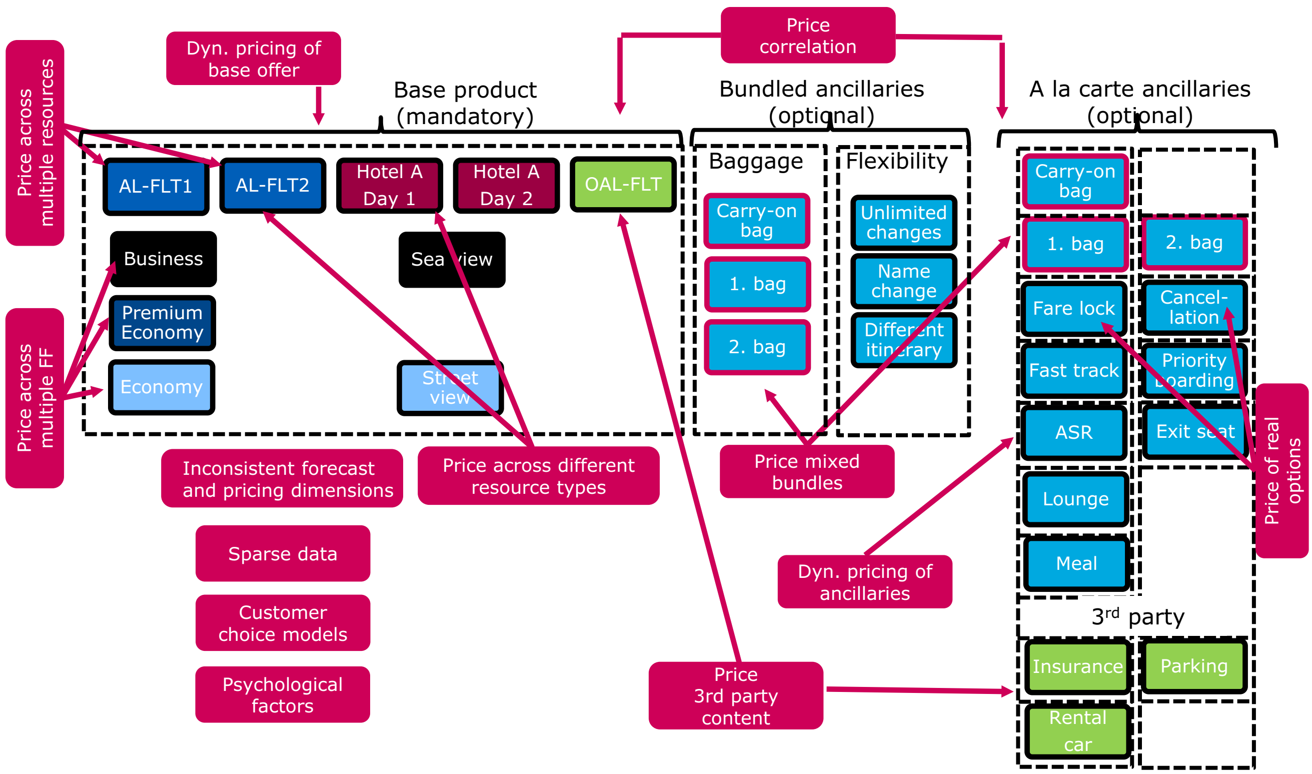

Fig. 1 Pricing challenges when pricing a general product offer. Red boxes indicate scientific issues to be overcome. Offer items indicated in dashed boxes. Offer items can have mandatory or optimal statuses

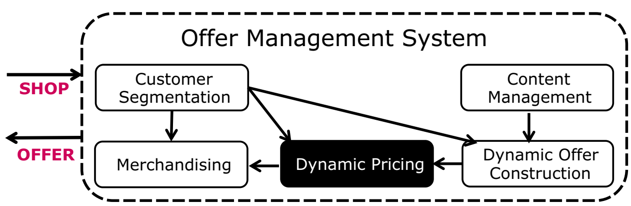

Fig. 2 Offer Management System (OMS). Shopping requests are passed to the OMS. Offers are returned to the customer

explained above. We assume that the catalog is predefined and that during the shopping flow, a product offer is constructed dynamically by selecting a set of products from the catalog.

Customer segmentation The underlying assumption behind customer segmentation is that we need to understand the customer's behavior to construct appropriate offers. This requires that existing and future customers can be categorized into customer segments, based on such shared characteristics as common needs, interests, lifestyles, or socio-demographics, past purchases, to name a few. Customer segmentation is performed offline, shared with other key functions in the airline, such as the digital user experience or servicing. A segmentation module must also consider cases when very little is known about the customer's identity (e.g., an anonymous search), deducing the customer's travel purpose from the search context. The calibrated customer segmentation model is loaded into OMS for online execution.

Dynamic offer construction This module recommends the most relevant set of products in order to maximize the conversion of the shopping request at an individualized/customer segment level. Offer construction can be optimized by recommender systems, which find broad application outside the airline industry (refer to the pioneering work by Linden (2003) or the recent textbook by Aggarwal (2016) for an overview of the field). The calibration of the models is performed offline and subsequently loaded into OMS for online execution. The module has several applications:

- *Rank products from the product catalog*, according to relevance.
- *Bundle product sets* to control the choice set available to the customer.
- *Propose upsell, cross-sell, and ancillaries* within shopping or re-shopping flows.
- *Send unsolicited notifications* (e.g., e-mail, mobile, or check-in kiosk) to customers to propose ancillary products.

Dynamic pricing The DP module prices the product offer that was constructed by the previous module. We will detail this module in Sect. "Dynamic pricing of offers."

Merchandizing From an RMS perspective, the airline is indifferent as to whether a customer accepts or rejects an offer. The assumption is that once a fair market price is set, the airline is better off rejecting customers who will not pay the price because other customers will. The reasoning being that this sort of all-or-nothing course overall results in higher revenue for the airline. However, in reality, the customer's purchasing behavior is neither rational nor guided by utility maximization. Customers are humans who can be affected by merchandising techniques.

The purpose of the merchandizing module is to apply merchandizing techniques, such as *framing* (the way offers are presented), *priming* (influencing by focusing on specific attributes), *defaults* (preset options), *decoys* (adding inferior offers), or *positioning* (specific position in the ranking) to affect the customer's purchasing behavior (Jannach et al. 2010).

Scientific challenges of pricing an offer

Since the 1980s, airlines have relied on RMS to price flight products. Therefore, it is natural to consider whether RMS principles could be applied in OMS. Comparing the objective of OMS against that of RMS (*to maximize revenue by pricing flight products*), there are two noteworthy differences. First, while RMS optimizes only the price of the flight product (actually setting availabilities by booking class having pre-filed fares), OMS optimizes both *constructing* and *pricing* of the offer. This is a much more ambitious undertaking. Second, in contrast to RMS, which provides the same price to all customers having the same itinerary and fare product, OMS may differentiate both products and prices at an individualized level (so-called personalized pricing). This practice may raise legal and moral concerns, such as violations of perceived fairness that would have to be addressed (Reinartz 2002).

Numerous challenges arise in moving toward the objective of OMS. Solutions differ in terms of their maturity, which we group into three categories: *mature* (a satisfactory solution exists that has been implemented in RMS); *emergent* (new ideas or theories exist but solutions have not been implemented); and *immature* (challenges that remain unsolved despite much research effort).

Mature (solved)

a. *Single resources* The optimization of a single flight leg dates back to Littlewood (1972), who solved the optimization problem for two fare classes. Over the years, this solution has been extended to multiple classes (Belobaba 1989) and further to dynamic optimization (Lee and Hersh 1993; Lautenbacher and Stidham 1999).
b. *Multiple resources* Offers may consume multiple resources, such as a flight connection or a hotel stay with multiple nights. The offer, however, must be accepted or rejected an entity. This means we cannot independently optimize resources. This challenge was solved in the early 1990s with the invention of Network RMS (Smith and Penn 1988; Williamson 1992).
c. *Fare family fare products* The customer's reservation prices for different fare products are correlated. This means that prices cannot be determined for each fare product independently but must be determined simultaneously. This challenge was solved with the advent of fare adjustment theory (Fiig et al. 2010, 2011).
d. *Mix of different resource types* Offers may include a bundle with a mix of different types of resources, such as a vacation package that includes a flight product and a hotel room. To manage this complexity, RM control has to apply additive bid prices across the different resource types—a natural extension of bid-price control (Weber 2001).
e. *Third-party content pricing* This is not a trivial issue, owing to complex SPA agreements for interline connections. For non-flight ancillary products, such as hotels, insurance, or parking, various commercial models may apply either an agency model or a merchant model, each of which has unique revenue maximization implications. However, the NDC standard proposes that an interline product is created as a contract between the offer responsible airline and its interline partner at the time of shopping (Hoyles 2015). From the perspective of the offer responsible airline, third parties can then be viewed as subcontractors and their prices treated as variable costs.

Emergent (new ideas)

f. *Dynamic pricing of flight product* The purpose of DP is to maximize contribution by dynamically pricing flight products, considering both customer and contextual information. In this regard, we can think of DP as a light version of OMS, with a narrow scope limited to flight products. Simulation studies by Fiig et al. (2015) have demonstrated that DP, by optimizing the contribution within the shopping session, can significantly improve revenue performance.
g. *Sparse-data issues in demand forecasting* In network RMS, demand forecasting is performed at the traffic

flow/point-of-sale/fare product dimensions. At this level of detail, thick flows ($\geq$ 10 passengers/year) account for only about 5% of all flows (Fiig 2007). Even when only thick flows are considered, data sparseness poses significant challenges (Gorin 2012; Fiig et al. 2014; Rauch et al. 2015, 2017). In OMS, the airline's offer is expanded to contain multiple products—creating a need for an even finer level of detail and thereby aggravating the data-sparseness issue.

h. *Inconsistent forecast and pricing dimensions* For the airline's pricing department, the dimensions provided by demand forecasting are not detailed enough. For this reason, the pricing department utilizes its own dimensions by segmenting customers by eligibility and fare rules. The inconsistency between forecasting and pricing dimensions lead to erroneous demand and WTP forecasts. To resolve this issue, as well as the data-sparseness issue, Rauch et al. (2015, 2017) proposed a groundbreaking idea of disentangling capacity control from price optimization. We will return to this concept in the next section.
i. *Customer choice models* To model individual customers' decision-making processes, discrete choice models have been found to be superior to traditional forecasting methodologies (Garrow 2010). Setting parameters for discrete choice models; however, requires information not only about the choice selected (as with traditional forecasting models) but also about the alternatives not chosen. This information has traditionally been difficult to obtain for the airlines, but technological advancements have made this more accessible now.
j. *Dynamic pricing of ancillary products* Today, airlines rely on static prices, filed in advance, for ancillary products. Prices may differ by market, time to departure, touch point, or sales channel but are not optimized. As such, airlines miss a revenue opportunity. In this paper, we show how ancillary products may be dynamically priced in each shopping session using price testing in conjunction with machine learning (see Sect. "Dynamic pricing of ancillaries").

Immature (unsolved)

k. *Mixed bundles* Products can be sold as pure products (a la carte), bundles (sold as an entity), or a combination of both (mixed bundles). This requires internal pricing consistency, because the customer can select from among multiple offers containing the same (or similar) products. For example, the carry-on bag and check-in bags, sold both a la carte and as a bundle in Fig. 1, offers six options for the customer (no bags, only carry-on bag, one check-in bag, two check-in bags, carry-on bag and one check-in bag, baggage bundle). The concepts were formulated in a seminal paper by Adams and Yellen (1976)—yet despite decades of research, no practical solution has been found.
l. *Correlated reservation prices* Consumers' reservation prices for different products may be correlated. This means that prices cannot be determined for each product independently but prices for all correlated products must be determined simultaneously. This issue already exists in RMS in fare families and is compounded when ancillaries are added because of the multitude of offers that can be constructed. Bockelie (2017) recently proposed a new dynamic programing approach that explicitly accounts for ancillary products and passenger choice behavior in RMS.
m. *Pricing of ticket options* Airlines may offer derivative products, such as ticket options (also called "time to think" or "fare locks") that are equivalent to a call option. This gives the customer the right—but not the obligation—to purchase the ticket within a given window of time (such as 1 week) at a given price. Pricing ticket options, however, is significantly different than pricing financial options. We will not go into detail here but refer the reader to our forthcoming paper (Sahin forthcoming).
n. *Psychological factors* The merchandizing techniques described above—framing, priming, defaults, decoy, positioning—influence human behavior. Much research [notably, Kahneman and Tversky's development of prospect theory (1979)] has been performed on behavioral economics. However, psychological factors so far have been difficult to model and consider in a rigorous scientific way.

Dynamic pricing of offers

We will assume that a product offer has been constructed and is ready to be priced. For this purpose, we need a precise definition of the objective of dynamic pricing:

Maximize contribution[2] by dynamically pricing the product offer, considering both customer and contextual information.

[2] *Contribution* is revenue less variable costs. Revenue Management often ignores the variable cost because the incremental cost of flying one additional passenger was small. However, with lower ticket prices and increasing fuel prices, landing fees, and taxes, this is no longer true. Further, for third-party content, the cost component cannot be ignored.

Assumptions

To progress, we need to introduce some simplifying assumptions. We will assume the notion of a fair market price, meaning that all customers in a given segment receive the same price for the same product offer. We will ignore the complexities of *mixed bundles*, *correlated reservation prices,* and *psychological factors* mentioned above. In the Sect. "Future research directions," we will discuss how we may relax these assumptions.

With these simplifying assumptions, the task of dynamic pricing is essentially reduced to dynamically pricing of independent offer items: the airline base product and the a la carte ancillaries. This is described in the next two sections.

Dynamic pricing of the airline base product/flight component

When maximizing the contribution for a given shopping request, we are evaluating the difference between the contribution of selling the product now and the contribution of *not* selling the product (opportunity cost). This evaluation is made complex by the fact that the opportunity cost for an airline's base product is non-zero, arising from demand pressure on the flight resources.

For this reason, demand forecasts should be detailed enough—typically, at the levels of traffic flow, fare class, point of sale, departure day, and days to departure—to allow airlines to control their flight resources. Even though this level of detail gives rise to significant data-sparseness issues, it still is not detailed enough for pricing. Therefore, the airline's pricing department utilizes their own dimensions, independent of RMS, on the basis of customer dimensions such as customer eligibility, minimum/maximum stay duration, stopover, combinability, round-trip restrictions, Saturday/Sunday restrictions, and advance-purchase restrictions, rules and restrictions—that differentiate customers' WTP. This inconsistency results in RMS that cannot accurately estimate or represent WTP along the flight resource dimensions.

To resolve these issues, Rauch et al. (2015, 2017) proposed a groundbreaking idea: disentangle capacity control from price optimization. Unlike traditional RMS, where a single forecast is used to forecast both demand (input for the optimizer that calculate bid prices [BP]) and WTP (input to the marginal revenue transformation that calculate fare modifiers [FM]), Rauch et al. propose splitting the forecast into *two completely independent* forecasts. A *price-elasticity* forecast used for *price optimization* (Fig. 3, left box), and a *demand* forecast used for *capacity control* (Fig. 3, right box).

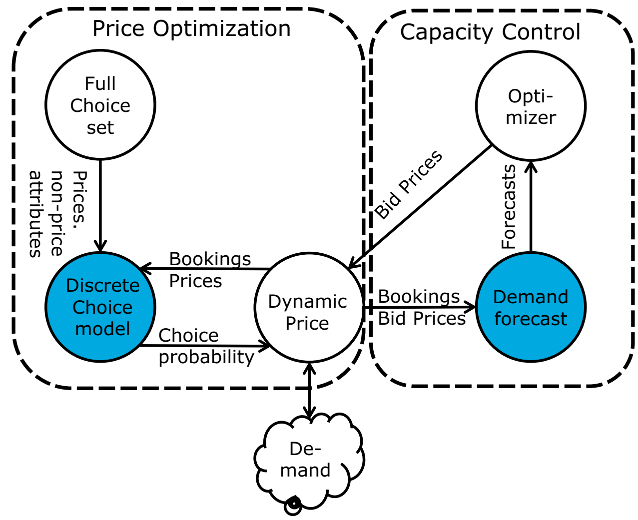

Fig. 3 Disentanglement applied to dynamic pricing

Based on this work, we now generalize the concept of disentanglement to enable *Dynamic Pricing*—one of the objectives of OMS. Capacity control is managed by a demand forecast that determines volume and network contribution, which through optimization produces bid prices (as originally proposed by Rauch et al.). To achieve price optimization, the price-elasticity forecasts must be extended to discrete choice models that consider the customer's full choice set (during the shopping request)—not just different booking classes of the traffic flow, as considered by RMS. In addition, the availability computation considered in RMS should be extended to take the general choice probabilities into account (see below).

Calibration of choice models

To model an individual customer's decision-making process, we apply discrete choice models. The data set for calibrating the choice models is constructed by matching bookings (from MIDT and PNRs from the airline's reservation system) with the corresponding shopping context—a comprehensive list of alternatives available to the customer at time of booking. This provides a data set in which we have, for each booked customer n, information about the alternatives i in their choice set $i \in C_n$, including prices p_i and non-price attributes $\boldsymbol{x}_{ni}$ (vector notation). Typical non-price attributes could be airline preference, time of day preference, and schedule quality (e.g., number of stops or trip duration).

In practice, not all bookings can be traced and set into their shopping context. In addition, the choice set may be incomplete, because GDS-based search transactions do not always capture information from all airlines (some airlines do not or only partially participate in GDS channels); nor do they capture information from other modes of transportation (e.g., rail or bus) that may be available to the customer.

Despite these limitations, however, we are now in a position to calibrate the selected choice model to obtain parameter estimates $\boldsymbol{\beta}$ and the choice probability $P(i|p_i, \boldsymbol{x}_{ni}, \boldsymbol{\beta})$. The choice model parameters $\boldsymbol{\beta}$ typically are linearly coupled to the attributes through the utility and as such interpreted as sensitivity parameters.

In practice, the models that have attracted the most attention are multinomial logit (MNL) and nested logit (NL) models. These models lead to a closed-form expression for the choice probability, which greatly simplifies both interpretation and computation.

Dynamic pricing

The objective of DP is to determine the prices that maximize the contribution of the offer. Consider the airline's base product only, for customer n, with a given itinerary $i \in C_n$. We can determine the optimal price p_i^* by

$$p_i^* = \underset{p_i}{\mathrm{Argmax}}\ [(p_i - BP)P(i|p_i, \boldsymbol{x}_{ni}, \boldsymbol{\beta})],$$

where BP denotes the origin-and-destination (O&D) bid price, obtained as the output of the capacity control in Fig. 3. It should be noted that, unlike the filed fares in RMS, the optimal price determined by DP is a continuous value.

Further, if we reduce the choice set, as assumed by RMS, to include only the different booking classes of a given traffic flow and limit ourselves to the discrete pre-filed fares f, the availability control decision in RMS is recovered. To see this, consider for simplicity a fenceless fare structure with an exponential sell-up probability $P^{RMS}(f) = e^{-\beta f}$, as seen by RMS. By inserting $P^{RMS}(f)$ as the choice probability in the DP equation above, the optimal RMS fare f^* becomes:

$$f^* = \underset{f}{\mathrm{Argmax}}\ \left[(f - BP)P^{RMS}(f)\right] = BP + 1/\beta,$$

where the term $1/\beta$ can be identified as the fare modifier. This expression is identical to RMS' bid-price acceptance criterion: accept request with $f \geq f^*$ and reject otherwise.

Dynamic pricing of ancillaries

For discussion purposes, we will assume that there is an unlimited supply of a la carte ancillary products. This assumption is often justified, although not always (there is limited capacity of exit seats or cabin overhead space).

Our objective is to maximize the contribution for a given customer request, which, as before, is done by evaluating the difference between the contribution of selling the ancillary now and the contribution of *not* selling the ancillary (the opportunity cost). Unlike before, however, the opportunity cost for ancillaries is zero because of our assumptions of unlimited supply (and uncorrelated reservation prices). This means we no longer have to perform demand forecasting (or optimization). Without this constraint, we are no longer limited to the traditional forecast dimensions tied to the flight resources and are free to use any set of attributes. Below, we provide a list, by no means exhaustive, of possible attributes that could be used:

- *Flight resource dimensions*: O&D market, departure day and time, days to departure, fare product.
- *Point of sale* (*POS*).
- *Booking information*: booking day, weekday, booking hour.
- *Trip information*: total travel time, number of connections, return travel date.
- *Time zones*: origin, destination time zones.
- *Price*: average price paid per passenger for the flight product.
- *Persona*: customer segment, travel purpose.
- *Number of passengers in party* (*NIP*).
- *Weather*: season, weather forecast at destination.
- *Holidays*: days to Thanksgiving, Christmas, Easter, etc.

Letting $\boldsymbol{x}_n$ denote the attributes of the customer request (e.g., among those in the list above) and $P(p, \boldsymbol{x}_n, \boldsymbol{\beta})$ denote the purchase probability for the considered ancillary at price p, then one approach could be to maximize directly (as before) the contribution: $p^* = \mathrm{Argmax}_p[p \cdot P(p, \boldsymbol{x}_n, \boldsymbol{\beta})]$. However, this would require us to calibrate (that is, estimate the $\boldsymbol{\beta}$ parameters) of $P(p, \boldsymbol{x}_n, \boldsymbol{\beta})$. A limitation of this approach is that a calibration of the purchase probability using traditional statistical methods requires low dimensionality of the attributes, such as $\dim(\boldsymbol{x}_n) \leq 5$, in order to collect sufficient data to perform meaningful statistical analysis.

Instead, we can directly apply A/B testing in conjunction with Machine Learning (ML). These techniques are able to handle high-dimensional data and can extract the significant pricing attributes without having to model customer choice. Typically, in our tests, the resulting significant attributes have $\dim(\boldsymbol{x}_n) \approx 10$.

In the remaining part of this section, we explain the principles behind Navitaire's product, Ancillary Price Optimization (APO) that applies this methodology.

Several low-cost and hybrid carriers across Europe, Asia, and North and South America use APO in production to experiment with dynamic pricing of ancillaries. The airline usually starts with bags or seats because of the higher sales volume and conversion rates that allow us to obtain the required data more quickly. But the principles apply generally, and so APO is also used to price other ancillary products, such as fast-track security, and bundles of ancillaries.

APO is conducted in five sequential steps (Coverston 2016), which we will illustrate by way of example.

1. *Configure the experiments* Defines the ancillary product; the target population (channel, market, customer segment, etc.); the test duration (e.g., 3 months); the split between traffic getting control price versus traffic for price testing; and specifications of the price point to be tested. In this example, 92% of traffic receives the control price (no treatment, price point 2), while the remaining 8% is randomly split between price point 1 or price point 3 (Fig. 4). The target percentage selected for price testing (in this example, 8%), is based on a projection of the time required to gather the data. The projection depends on conversion rates for flights and the ancillaries.
2. *Run the experiment* The data collected from this experiment include the attributes of the request, price point, and outcome (purchase/no purchase). Summary statistics for the pricing experiment in this example are shown at the bottom of Fig. 4. The numbers above the price point denote the probability $P(p_i)$ of the treatment, while numbers below denote the purchase probability $P(\text{purchase}|p_i)$ and the expected revenue $r_i = P(\text{purchase}|p_i)p_i$, given the treatment, respectively. It can be seen that the price point that provides the highest expected revenue is p_2.
3. *Analyze and model* Conduct an offline analysis and model selection based on data gathered during experimentation. Significant pricing attributes are identified and a model is selected among standard methods (e.g., regression, decision trees, random forest, neural networks) based on its implications for marketing and brand in addition to the revenue uplift (see below). Figure 5 shows the optimized model with three decision trees—one for each price point, although

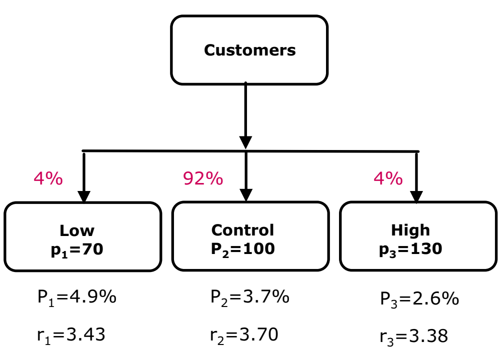

Fig. 4 Randomized price experiment. The numbers above the price point denote the probability $P(p_i)$ of the treatment, while numbers below denote the purchase probability $P(\text{purchase}|p_i)$ and the expected revenue, $r_i = P(\text{purchase}|p_i)p_i$, given the treatment

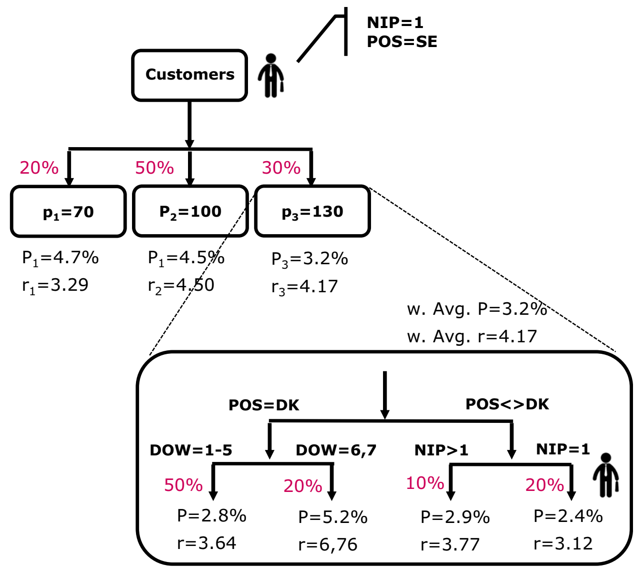

Fig. 5 Optimized model, applied in real time. Point of sale (POS) and number in party (NIP) are shown as attributes in the decision tree. The numbers above the price point denote the probability $P(p_i)$ of the treatment, while numbers below denote the purchase probability $P(\text{purchase}|p_i)$ and the expected revenue $r_i = P(p_i)p_i$, given the treatment

only the decision tree for p_3 is shown. Summary statistics can be computed to assess the performance of the optimized model. From the decision tree, we can compute the average purchase probability by the weighted average over the purchase probabilities of the terminal nodes. For the decision tree p_3, we have average purchase probability $P = 3.2\%$ and average expected revenue $r = 4.17$.
4. *Deploy pricing model* Provided that the uplift analysis (see below) shows sufficient benefit, the model is deployed to production. Continued experimentation is normally preferred; for instance, our example in Fig. 4 could be extended such that 92% of traffic is again allocated to the control group (using the existing model) while the remaining 8% is directed toward new experimentation. Figure 5 illustrates real-time execution, using our example. Consider a customer; traveling alone (NIP = 1), with point-of-sale Sweden (POS = SE). If we propose him with the price point $p_3 = 130$, he would end up with expected revenue $r = 3.12$. Similar calculations are performed for price points, p_1 and p_2. The expected revenue of all three price points are compared and the price point with the highest expected revenue is proposed to the customer.
5. *Monitor* This step provides statistics to verify that the optimized pricing model behaves according to expectations. The cycle is repeated on a continuous basis as new pricing experiments continue even after the optimized model is deployed.

Expected uplift

To assess the effect of the pricing experiment, we apply uplift models that quantify this effect by measuring the difference in conditional purchase probability between the treatment and non-treatment (control) customers (Radcliffe 2007). However, depending on the objective (e.g., revenue maximization), the conditional probabilities will in general be weighted. In our example case, we would like to compute the expected revenue uplift:

$$u = \left[\frac{\sum_{i=1}^{3} p_i \cdot P(\text{purchase}|p_i)P(p_i) - p_c \cdot P(\text{purchase}|p_c)}{p_c \cdot P(\text{purchase}|p_c)}\right]$$

To provide the most conservative uplift estimate, we choose as a control the price point that provides the highest revenue among the three price points. In the example case, (Fig. 4), this is obtained for $p_{\mathrm{c}} = p_2$, with revenue $r_2 = 3.7$. Performing the uplift computation yields

$$u = \frac{20\% \cdot 3.29 + 50\% \cdot 4.5 + 30\% \cdot 4.17 - 3.7}{3.7} = 12.4\%$$

While we have illustrated the methodology for an example case, the uplift is of correct order of magnitude for real airline applications. Hence, this approach provides substantial revenue potential versus the current practice of static pricing.

A distribution landscape in motion

Looking ahead, one of the biggest challenges for deploying dynamic pricing of airline offers is the management of volume and the complexity of worldwide travel distribution. Today, GDSs are handling high volumes of shopping transactions with multi-airline aggregation, though for simpler airline offers and with more homogeneous pricing and fare definitions. In this section, we will briefly describe current industry initiatives, then review characteristics of the distribution channels and their impact on airlines' OMS.

Industry initiatives enabling dynamic pricing

The design of a new distribution standard, known as IATA NDC, was initiated several years ago. This standard is intended to facilitate the connection between third-party retailers and an airline's offer and order management systems, with a goal of avoiding prohibitive bilateral integration costs between each airline and third-party distribution channel.

More importantly, the standards also seek to give airlines full control of the offer generation process, including access to persona and contextual information and ability to dynamically construct and price offers. With NDC, airlines avoid the two-step process of fare filing, and availability computation and distribution, which has prevailed for decades (Hoyles 2015; Wilson and Touraine 2016).

Figure 6 illustrates where the airline OMS fits in this new distribution landscape.

According to IATA (IATA.org), there are now more than 40 airlines NDC-capable with a certification for offer and order management (Level 3); more than 30 companies providing those IT capabilities to airlines; and already a few aggregators and resellers capable of managing offers and orders using these new standards. Most airlines currently limit the deployment to direct connect to specific distribution partners, such as corporate customers, online travel agencies (OTAs), travel management companies (TMC), or other travel agencies (TAs).

Few leading airlines are now taking first steps toward a multi-channel retailing strategy focusing on an extended offering (IATA 2018). However, there is no known deployment of dynamic pricing for flight products using NDC to date.

In parallel, ATPCO has initiated a working group (Ratliff, 2017) focusing on dynamic pricing engine (DPE) deployment, which would re-use the legacy GDSs and airline CRSs pricing engines adding price adjustments. The group's objective is to reduce time to market compared with a full NDC and OMS deployment. Pilot testing to refine the design of a DPE is in progress with providers.

Distribution channels and their interactions with OMS

Let us now describe the different distribution channels and their main characteristics that could impact an airline's OMS.

The airlines' own digital touchpoints (e.g., website, mobile application) are critical to the customer experience and loyalty; those are already *directly connected* to the airline's OMS via enriched application programing interfaces (API), allowing for innovation and personalization of the offer. By controlling how offers are presented (e.g., through graphics, photography, or descriptions), airlines can also apply retailing techniques to increase conversion and upsell.

In reality, before reaching an airline's website, most customers spend time on the Internet's giants web interfaces, social media, and search engines (GAFA, shorthand for Google–Amazon–Facebook–Apple; and BAT, shorthand for Baidu–Alibaba–Tencent) to start their travel inspiration journey. To reach potential customers and create awareness about their offers, airlines need solutions for distributing their offers through these *gatekeepers*, to

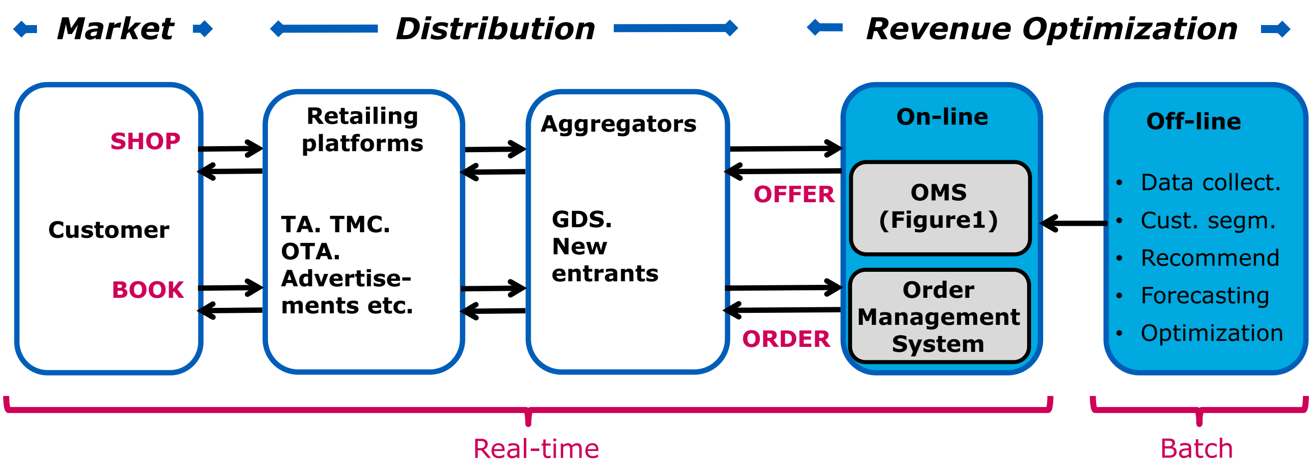

Fig. 6 Customer shopping request in an agent's front-office system is passed through the aggregator and to the airline's OMS. The airline creates the offer that is returned to the customer. Each offer is individually tagged with an offer ID that subsequently is used if the customer accepts the offer ("book"), in which case an order is created

which airlines pay advertising or referral fees. This customer acquisition step is not the same as a standard shopping step when a customer is actively searching for a travel product. As a consequence, these digital marketing channels require specific solutions coming from the airline OMS, to construct and promote the most relevant offers.

Meta-searchers (such as Kayak, Momondo, and Cheapflights) have a specific role in customer acquisition, as they answer a customer need to compare offers among all travel retailers. Meta-searchers generate a huge volume of search traffic that requires an instant answer, putting significant stress on airlines and aggregators' IT systems and without any guarantee of the customer completing a booking. As the customer's conversion is sometimes impacted by redirection to another website, meta-searchers also propose including a "meta-booking" capability in their shopping tools, while directly connecting to retailers in the back for order fulfillment (for instance, through NDC).

Third-party travel retailers (TAs, OTAs, and TMCs) account for the majority of bookings on full-service carriers. These retailers also need aggregated information from all airlines to offer a variety of choices to their customers. While these retailers typically rely on the GDSs to provide this service, they can also implement a direct connection to some airlines, for instance through the NDC standard, as described earlier. Of note, they also have their own advanced retailing techniques, thus controlling which itineraries to display and in which order, based on commissions, performance-specific incentives, and service fees (Smith et al. 2007).

Technological advancement

As online search engines generate thousands of search transactions per booking, an IT system needs to be extremely cost-effective, scalable, provide real-time dynamic offer construction and dynamic pricing while providing consistency across all distribution channels. These needs have to align with the airline's and distribution partners business priorities and the expected return on investment.

Technology will help to accomplish these needs. Cloud infrastructure and real-time worldwide data synchronization allow data centers across continents to host and run a single source of content, accessible to any distribution channel, while continuously being under airline control.

How does this work? The IT system managing the airline's OMS (including algorithms and data), is replicated in real time on a platform that uses a virtualized infrastructure. This platform can be deployed in different data centers, providing unprecedented scalability and resiliency while improving network response time because of the physical proximity to the airlines' consumers. Further, cloud infrastructure providers make significant investments in security, automation, and cost control, providing the tools necessary to optimize performance and processing costs.

Future research directions

While many pricing challenges, as mentioned in this paper, have so far proven intractable from a traditional modeling perspective, the new paradigm in science is ML. Analogous to the methodology explained for Navitaire's APO, we may apply price experimentation in the production system in conjunction with ML to optimize relevant key performance indicators, such as revenue uplift or conversion rates.

As direction for future research, we propose applying ML techniques to address the pricing complexities of mixed bundles, correlated reservation prices, and psychological factors.

Conclusion

We have presented our vision for an airline offer management system that enables airlines to dynamically construct and dynamically price offers. Fulfilling this vision is not possible today, given the limitations of legacy RM and distribution systems.

We explain why existing RM techniques are inadequate for pricing general offers and provide insight into new emergent techniques. We explain how these techniques can be employed to perform dynamic pricing of both the airline's base product and ancillary products.

In light of NDC, we discuss the transformation of the distribution landscape and its implications. The digital revolution that has penetrated many industries will also affect travel distribution. New entrants, Internet giants, and meta-searchers change customers' expectations to instant, relevant, and personalized offers—delivered consistently across all channels and devices. Finally, we discuss how technological advancements in cloud infrastructure and real-time worldwide data synchronization can enable deployment of Dynamic Pricing of the airline offers.

Acknowledgements We thank Jason Coverston, at Navitaire, for sharing his insights on Navitaire's Ancillary Price Optimization (APO). We also thank Stephane Lecourtois for sharing his knowledge on merchandising.

References

Adams, W.J., and J.L. Yellen. 1976. Commodity Bundling and the Burden of Monopoly. *Quarterly Journal of Economics* 90 (3): 475–498.

Aggarwal, C.C. 2016. *An Introduction to Recommender Systems*. New York: Springer.

Belobaba, P.P. 1989. Application of a Probabilistic Decision Model to Airline Seat Inventory Control. *Operations Research* 37: 183–197.

Bockelie, A. 2017. Revenue Management with Ancillary Services: A New Optimization Approach. Presented at AGIFORS Reservation and Distribution Study Group, Frankfurt, Germany.

Coverston, J. 2016. Ancillary Pricing Optimization (APO), Navitaire Customer Conference, Las Vegas.

Hoyles, Y. 2015. IATA Simplifying the Business. New Distribution Capability (NDC). Facilitating Air Retailing. http://www.iata.org/whatwedo/airline-distribution/ndc/Documents/ndc-strategy-paper.pdf. Accessed 15 Oct 2017.

Fiig, T. 2007. ODYSSEY: SAS' O&D Forecasting System. Presented at AGIFORS Reservation and Yield Management Study Group, Jeju Island, Korea.

Fiig, T., K. Isler, C. Hopperstad, and P. Belobaba. 2010. Optimization of Mixed Fare Structures: Theory and Applications. *Journal of Revenue and Pricing Management* 9 (1/2): 152–170.

Fiig, T., K. Isler, C. Hopperstad, and S. Olsen. 2011. Forecasting and Optimization of Fare Families. *Journal of Revenue and Pricing Management* 11 (3): 322–342.

Fiig, T., B. Smith, G. Goyons, and R. Adelving. 2015. Dynamic Pricing—the Next Revolution in RM? *Journal of Revenue and Pricing Management* 15 (5): 360–379.

Fiig, T., R. Härdling, S. Pölt, and C. Hopperstad. 2014. Demand Forecasting and Measuring Forecast Accuracy in General Fare Structures. *Management* 13 (6): 413–439.

Garrow, L. 2010. *Discrete Choice Modeling and Air Travel Demand: Theory and Applications*. Farnham: Ashgate Publishing.

Gorin, T. 2012. Forecasting Passenger Demand at the "Right" Level. Optimal Dimension Definitions and Clusters for Forecasting Demand. Presented at AGIFORS Reservation and Distribution Study Group, Barcelona, Spain.

IATA. 2013. Profitability and the Air Transport Value Chain. IATA Economics Briefing No. 10, IATA.

IATA. 2016. New Distribution Capability (NDC) Implementation Guide 3.2. http://www.iata.org/whatwedo/airline-distribution/ndc/Documents/ndc-implementation-guide.pdf. Accessed 3 Mar 2018.

IATA. 2018. NDC Airlines. http://www.iata.org/whatwedo/airline-distribution/ndc/Pages/ndc-airlines.aspx. Accessed 15 Feb 2018.

Jannach, D., M. Zanker, A. Felfernig, et al. 2010. *Recommender Systems: An Introduction (Chapter 10)*. Cambridge: Cambridge University Press.

Kahneman, D., and F.A. Tversky. 1979. Prospect Theory: An Analysis of Decision Under risk. *Econometrica* 47 (2): 263–292.

Lautenbacher, C.J., and S.J. Stidham. 1999. The Underlying Markov Decision Process in the Single-Leg Airline Yield Management Problem. *Transportation Science* 33: 136–146.

Lee, T.C., and M. Hirsch. 1993. A Model for Dynamic Airline Seat Inventory Control with Multiple Seat Bookings. *Transportation Science* 27: 252–265.

Linden, B., J. Smith, and J. York. 2003. Amazon.com Recommendations: Item-to-Item Collaborative Filtering. *IEEE Internet Computing* 7 (1): 76–80.

Littlewood, K. 1972. Forecasting and Control of Passenger Bookings. 12th AGIFORS Annual Symposium Proceedings, Nathanya, Israel, pp. 95–117.

Popp, C. 2016. RM Meets Distribution. Presented at AGIFORS Reservation and Distribution Study Group, San Francisco, CA

Radcliffe, N.J. 2007. Using Control Groups to Target on Predicted Lift: Building and Assessing Uplift Model. *Direct Marketing Analytics Journal* 1: 14–21.

Ratliff, R. 2017. Industry-standard Specifications for Air Dynamic Pricing Engines: Progress Update. Presented at AGIFORS Revenue Management Study Group Meeting, San Francisco, USA.

Rauch, J. Pölt, S., Isler, K. 2015. Disentangling Capacity Control from Price Optimization. Presented at AGIFORS Reservation and Distribution Study Group, Shanghai, China

Rauch, J., K. Isler, and S. Poelt. 2017. Disentangling Capacity Control from Price Optimization. *Journal of Revenue and Pricing Management*. https://doi.org/10.1057/s41272-017-0118-9.

Reinartz, W. 2002. Customizing Prices in Online Markets. *Symphonya Emerging Issues in Management* 1: 55–65.

Sahin, O. Hagen, J., Jadanza, R., and Fiig T. forthcoming. Dynamically Pricing Ticket Options: An Artificial Intelligence Approach.

Smith, B.C., and C.W. Penn. 1988. Analysis of Alternative Origin-Destination Control Strategies. *AGIFORS Symposium Proceedings* 28: 123–144.

Smith, B., R. Darrow, J. Elieson, D. Guenther, V. Rao, and F. Zouaoui. 2007. Travelocity Becomes a Travel Retailer. *Interfaces* 37 (1): 68–81.

Weber, K. 2001. From O&D Bid Price Control to Package Bid Price Control. Presented at AGIFORS Reservation and Yield Management Study Group, Bangkok, Thailand.

Williamson, E.L. 1992. Airline Network Seat Inventory Control: Methodologies and Revenue Impacts. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Wilson, B. and Touraine, S. (2016). NDC—ONE Order initiatives: Industry opportunities for simplification and value creation. Presented at AGIFORS Reservation and Distribution Study Group, Frankfurt, Germany.

Thomas Fiig is the Director, Chief Scientist at Amadeus, where he is responsible for revenue management strategy and scientific methodologies. He holds a Ph.D. in theoretical physics and mathematics and a B.A. in finance from the University of Copenhagen, Denmark. He has published several articles, recently focused on forecasting and optimization of simplified fare structures and dynamic pricing. He serves on the editorial board of the *Journal of Revenue and Pricing Management.*

Remy Le Guen leads the product management group for Revenue Management solutions at Amadeus. He has prior experience in management consulting, in the Revenue Management, Pricing and Revenue Integrity from Air France. Remy holds a M.Sc. in electrical engineering from Supelec, Paris, France.

Mathilde Gauchet is managing Amadeus Airline IT portfolio strategy. She headed Revenue Optimization and Inventory commercial solutions at Amadeus from 2009 to 2016, after contributing to the design and launch of the Amadeus Passenger Servicing System as part of the Research and Development. She holds an engineering degree from ENSTA, Paris, France and M.Sc. in Aeronautics from Caltech, Pasadena, CA, USA.