Dynamic offer generation in airline revenue management

Dynamic offer generation in airline revenue management

Cite this article as: KevinK. Wang, MichaelD. Wittman, Adam Bockelie, Dynamic offer generation in airline revenue management, Journal of Revenue and Pricing Management https://doi.org/10.1057/s41272-021-00349-4

Dynamic offer generation in airline revenue management

Kevin K. Wang*[a,1], Michael D. Wittman[b,2] and Adam Bockelie[c,3]

[*] Corresponding author

[a] Massachusetts Institute of Technology, International Center for Air Transportation, 77 Massachusetts Avenue, Building 35-217, Cambridge, MA 02139
Email: kevini@mit.edu

[b] Amadeus IT Group, Lufthavnsboulevarden 14, 2. tv, DK-2770 Kastrup, Denmark
Email: michael.wittman@amadeus.com

[c] Air Canada, 7373, bd de la Côte-Vertu Ouest, Saint-Laurent, QC, Canada, H4Y 1H4
Email: adam.bockelie@aircanada.ca

[1] is a Ph.D. student in Air Transportation Systems at MIT, where his research focuses on the optimization of ancillary services and dynamic offers in airline revenue management systems. He holds a S.M. in Aeronautics and Astronautics from MIT and an M.Eng. and B.A. in Engineering from the University of Cambridge.

[2] is Expert, Revenue Management Science & Research at Amadeus, where he works on developing new models for forecasting and optimization in revenue management and dynamic pricing. He holds a Ph.D. in Air Transportation Systems from MIT and is the author of a dozen articles published in peer-reviewed academic journals. His work has won awards and recognition from AGIFORS and INFORMS.

[3] is Manager, Revenue Management Strategic Planning at Air Canada where he is responsible for developing the long-term revenue management roadmap and for integrating artificial intelligence into revenue management practice. He holds a Ph.D. in Transportation from MIT, where he studied passenger choice modeling, total revenue optimization, and dynamic offer creation.

1

Dynamic offer generation in airline revenue management

Kevin K. Wang, Michael D. Wittman and Adam Bockelie

Abstract

The New Distribution Capability and new retailing platforms will enable airlines to respond to shopping requests with bundled offers of flights and ancillary services, representing an evolution from traditional, flight-focused optimization. Assembling an attractive set of offers to display to customers therefore represents a new joint pricing and assortment optimization problem in airline revenue management.

In this paper, we introduce an initial optimization approach for the selection and pricing of *a la carte* and bundled flight and ancillary offers. First, we propose a customer choice model that captures the impact of ancillary bundles on flight itinerary choice. We then calculate prices for each offer from a continuous range of price points and display the offer set that maximizes expected revenue for a given customer segment. We illustrate the approach using a single-flight, single-ancillary base case and discuss extensions to more complex environments.

Tests in the Passenger Origin-Destination Simulator (PODS) show that dynamic offer generation (DOG) can increase net revenue when used by one or more airlines in a competitive network, assuming that the airlines are able to accurately segment incoming requests and estimate the average willingness-to-pay of each segment. We find that the majority of the revenue gains of DOG are due to competitive effects from the dynamic pricing of the flight component of the offer. The bundling mechanism of DOG is a secondary source of revenue gain that can be realized when customers take bundled ancillary services into account when choosing the flight.

Our results provide insight for practitioners that are implementing offer optimization systems and processes. For example, in line with previous literature on bundle pricing, we find that in transparent distribution channels an ancillary service should be bundled with the flight when the valuation for the ancillary is high or when its marginal cost of provision is low. We close by discussing the strategic and managerial implications of a move from traditional distribution strategies to a next-generation, offer-focused approach.

Keywords: dynamic offers, dynamic pricing, assortment optimization, airline revenue management, New Distribution Capability

INTRODUCTION

An *offer* is a set of products and services that is displayed to a customer along with an associated price. In the airline industry, the construction and pricing of attractive and profitable offers – consisting of flight itineraries, core components such as the flight cabin and in-flight service, and ancillary services such as checked bags and extra legroom – is the central responsibility of the airline's commercial organization.

In a typical shopping session, airlines often display offer components to customers sequentially. For example, a customer may be shown a list of flight itineraries, followed by a choice of branded fare options, and then finally a selection of optional ancillary services. It is only after these consecutive decisions that the customer is shown a total price for the airline's offer. Airlines do not typically display cohesive offers that bundle a flight itinerary with one or more ancillary services at a single price.

There are several reasons why airlines have not moved to fully offer-centric retailing. First, current distribution and reservation system technologies limit the ways in which airlines can construct, display, and service bundled offers. For example, messaging standards limit what information can be passed between the customer and the airline on the indirect distribution channel. Even on the airline's own direct channel, technological limitations often restrict the degree to which airlines can sell and fulfill bundled offers.

Furthermore, silos in airline commercial organizations often lead airlines to optimize each component of the offer separately. For example, the schedule and network planning departments define the set of possible itineraries, the pricing and revenue management (RM) departments determine the fare products and fare class availability, and the marketing department designs the loyalty program benefits and prices the ancillary services. When constructing offers, airlines must also comply with legislative guidelines on the use of customer data and tariff filing requirements, which can be quite complex and vary by country.

Despite these challenges, recent innovations in distribution technology have led to new possibilities for offer-centric retailing. The International Air Transport Association (IATA)'s New Distribution Capability (NDC) allows products and services to be combined together into bundled offers that can then be distributed to customers at a single price (Westermann, 2013). NDC also allows for more customer-specific and context-specific information, such as frequent-flyer status and desired ancillary services, to be used during offer construction. This allows airlines to customize offers in real time to meet the needs of the customer making the request.

While NDC is still a new technology, the offer optimization it enables – in which flights and ancillary services are combined and priced in real time using context- or customer-specific information – is far more advanced than how airlines generate and price offers today. The scientific models that airlines currently use to optimize individual offer components are not fully suited for joint optimization of the entire offer. This paper aims to bridge this gap by introducing an approach for airline offer optimization that we call *dynamic offer generation* (DOG).

3

DOG selects and prices the optimal revenue-maximizing offer set, consisting of one or more offers, out of the full catalog of products that is offered by the airline. Offers, prices, and offer sets can be customized by customer type, allowing the possibility for contextualization based on the content of the search request or information about the customer. We show how this type of offer optimization can provide revenue benefits to airlines, due to the dynamic pricing of the flight and ancillary components of the offer and the bundling of ancillary services. We also assess the conditions in which it is favorable for airlines to offer only bundled offers versus offering ancillary services *a la carte*.

Unlike prior work, DOG is specifically designed for the airline offer optimization problem and incorporates information from a traditional airline revenue management system (RMS). Furthermore, our work is the first to simulate the impacts of offer optimization in a competitive airline environment, providing valuable insight to industry practitioners about the potential benefits and competitive effects of such a strategy compared to traditional RM.

The remainder of the paper is structured as follows: first, we discuss recent advancements in the offer optimization context in the airline industry literature and in other fields. Next, we formulate the general airline offer optimization problem for flights and ancillary services. We test our solution approach alongside a traditional RMS in a competitive simulation and show that dynamic offer generation can lead to revenue gains. We discuss the potential implications for industry practitioners, before concluding with recommendations for further work in this area.

LITERATURE REVIEW

Incorporating ancillary services into traditional airline RM

The airline RM literature has traditionally focused on the optimization of the flight seat alone, without incorporating optional ancillary services. The available prices for each flight itinerary are selected from a set of pre-priced filed fare classes. This setting, which we call *traditional RM*, is due to current distribution standards that limit the number of price points that airlines can make available in the marketplace.

Several studies have investigated how traditional RM could be extended to incorporate ancillary services. Hao (2014) proposed a simple heuristic to add the average revenue from the ancillary services for each fare class to the yields used by the RMS to optimize availability. Bockelie and Belobaba (2017) proposed a more complex customer choice framework to model how customers choose whether to purchase an ancillary service. They proposed two different customer behavior types: *sequential* customers, who make their flight itinerary selection unaware of the existence of ancillary services, and *simultaneous* customers, who incorporate their desired ancillary service into their flight itinerary selection process. Bockelie (2019) later designed a method based on the choice-based RM framework of Talluri and van Ryzin (2004) to incorporate this ancillary customer choice model into RMS forecasting and optimization.

These studies assume traditional pricing and distribution practices with filed fare classes, and do not focus on dynamic pricing of the flights or ancillary services. The prices of the ancillary services are not optimized in these studies.

Dynamic pricing of flight seats and ancillary services

Since NDC offers a possibility to distribute a wider range of price points (Dezelak and Ratliff, 2017), recent studies have also investigated dynamic pricing (DP) of flight seats and ancillary services. With DP, the price of the flight seat or ancillary service is selected from a continuous range of possible prices instead of from a set of pre-filed fares. A dynamic price could be represented either as an adjustment (increment or discount) from a filed fare, or without any reference to a filed fare (Wittman and Belobaba, 2019). DP also allows the use of contextual information about the shopping session – for instance, characteristics of the search request or the offers proposed by competitors – to further customize the price.

DP has been applied separately to the flight seat and to ancillary services. For the flight seat, several studies have shown that DP provides revenue benefits over traditional RM (Fiig et al., 2016; Kumar et al., 2017; Wittman and Belobaba, 2018). These benefits result from the customization of prices to each shopping session, as well as the increased range of price points that can be offered by DP. For ancillary services, Shukla et al. (2019) proposed a machine learning method for computing dynamic prices for one or more ancillary services based on past sales data, also considering customer context. They also found revenue benefits over static ancillary pricing models.

Joint optimization of flight seats and ancillary services

As NDC has become more mature, the literature has started to examine joint optimization and pricing of the flight seat and ancillary services. Several studies have discussed the scientific and technical components of the offer management systems (OMS) that would be needed to create, distribute, and fulfill offers (Madireddy et al., 2017; Fiig et al., 2018; Vinod et al., 2018). These studies serve as roadmaps for the design of such systems, but do not offer specific models for the pricing and optimization of bundled offers.

Several studies have proposed models for solving the joint optimization and pricing of flight seats and ancillary services. Ødegaard and Wilson (2016) model dynamic pricing of a flight and a single ancillary service for three types of ancillary purchase behavior types. They investigate in detail the pricing behavior assuming uniformly distributed willingness-to-pay. More recently, Shao and Kauermann (2020) estimate the price elasticities of airline customers facing *a la carte* or bundled offers using airline data. They investigate three different bundle pricing policies and estimate the change in expected revenue of shifting from one policy to another.

Joint optimization of bundled offers in the economics and operations research literature

While joint optimization of bundled offers is a relatively new concept in the airline industry literature, there are parallels to be found in the economics literature on bundle pricing and add-on pricing, and in the operations research on joint pricing and assortment

5

optimization problems. In the economics literature, Adams and Yellen (1976) provide a canonical model of the bundle pricing of two products. They show when it is best for the seller to bundle the products and when it is best to sell the products *a la carte*, depending on the valuations of the customers for the two products.

Ellison (2005), Gabaix and Laibson (2006), and Shulman and Geng (2013), among others, explore the problem of *add-on pricing*, in which the price of a core product is determined along with an optional add-on product that enhances the core product's quality. This can be compared to optional ancillary services that increase the quality of the travel experience, or to an airline fare family that bundles additional benefits with a base fare. However, these papers are highly theoretical in nature and do not offer specific heuristics or methods tailored for the airline offer optimization problem.

In the operations research literature, several studies have investigated *joint pricing and assortment optimization* where an optimal subset of products, perhaps subject to a display space constraint, is selected and priced to maximize the seller's expected revenue (Jagabathula and Rusmevichientong, 2017; Cataldo and Ferrer, 2017). These papers in the economics and OR literature do not typically consider a capacity constraint on the products contained in the offer, which is the case in the traditional airline RM problem.

In our paper, we focus on applying some of the general techniques of joint pricing and assortment optimization to the specific context of airline offer optimization. We extend the simultaneous/sequential customer choice framework first proposed by Bockelie and Belobaba (2017) to model customer choice amongst bundled offers and simulate the performance of our model compared to traditional RM and DP of the flight seat using the Passenger Origin-Destination Simulator (PODS).

OVERVIEW OF DYNAMIC OFFER GENERATION

The airline offer optimization problem

[Figure 1 about here]

Figure 1 shows a schematic of the airline offer optimization problem and its relationship to traditional pricing and RM. Airlines operate *flight itineraries* $f_j \in \{f_1, \ldots, f_J\}$ that each consist of one or more flight legs. Using a set of filed fares and a historical database of past booking behavior, a traditional RMS calculates flight leg bid prices as a function of a demand forecast and the remaining capacity on each flight leg. The flight itinerary bid prices π_j are computed by summing the bid prices for each flight leg in each itinerary f_j.

These flight leg bid prices are taken as an input to dynamic offer generation. Airlines also offer a catalog of *ancillary services*. Each ancillary service $a_k \in \{a_1, \ldots, a_K\}$ has a cost of provision m_k, which represents the cost to the airline of offering the service. For example, offering a customer lounge access may require the airline to make a payment to the lounge operator.

6

[Figure 2 about here]

In the Offer Set Generation process, airlines first construct *offers* $O_i \in \{O_1, \ldots, O_I\}$ from their catalogs of itineraries and ancillary services. An offer $O_i = \{f_j, a_1, \ldots, a_k\}$ contains exactly one flight itinerary and zero or more ancillary services. Each offer O_i has an associated *offer cost* $c_i = \pi_{j:f_j \in O_i} + \sum_{k:a_k \in O_i} m_k$. Since the flight itinerary bid price represents the cost to the airline of selling one seat on the itinerary, the offer cost represents the cost to the airline of providing the offer. Figure 2 shows an example where the airline has one flight itinerary and two ancillary services. In this case, there are four possible offers that the airline could display.

The airline then constructs an *offer set* $S = \{O_1, \ldots, O_n\}$ consisting of one or more offers. We define $\Omega = \{S_1, \ldots, S_M\}$ as the universe of all possible offer sets. Each offer O_i in an offer set has a corresponding *offer price* p_i. We define $\boldsymbol{p}(S) = (p_1, \ldots, p_n)$ as the vector of prices for all offers in offer set S, which is determined in the Offer Set Pricing process. The airline then selects a single offer set S^* to display to the customer in the Offer Set Selection process. In Figure 2, the displayed offer set contains offers O_2 and O_4.

Customer choice between offers

When faced with an offer set, a customer selects exactly one offer to purchase or chooses to purchase nothing. His choice depends on the other offers in the offer set, as well as the prices of the offers.

To model the probability that a customer selects a given offer from an offer set, we extend the continuous purchase behavior from the ancillary choice framework first proposed by Bockelie and Belobaba (2017). Customers draw a maximum willingness-to-pay (WTP) θ_j for the flight itinerary f_j and maximum WTPs γ_k for each ancillary service a_k. The maximum WTPs are drawn from independent random distributions, which could vary between different customer segments.

We assume that the customer's maximum WTP w_i for offer O_i is the sum of the offer's individual components and that the WTPs for the individual services are independent and uncorrelated:

$$w_i = \theta_{j:f_j \in O_i} + \sum_{k:a_k \in O_i} \gamma_k \qquad (1)$$

These assumptions are common in the bundle pricing literature (Adams and Yellen, 1976). Generally, the profitability of bundling has been shown to be greater when the WTP for bundled products are negatively correlated, meaning that a customer's valuation is typically lower for a second product if their valuation is high for the first product (Schmalensee, 1984). In the airline context, this can be seen with leisure-oriented ancillary services such as checked baggage: passengers willing to pay more for the flight are more likely to be business passengers with a lower valuation for baggage (we investigate such a relationship in our simulations). Among ancillary services, business-

7

oriented services such as fast track security could also be negatively correlated with checked baggage.

We assume that the customer chooses exactly one offer in the offer set, and that they choose the offer that has the highest utility (or consumer surplus) $u_i = w_i - p_i$, which is defined as the difference between the customer's maximum WTP for the offer and its price. Such *maximum utility models* are commonly used in the economics literature to model rational customer choice (Dobson and Kalish, 1988; Hanson and Martin, 1990).

The probability $Pr(O_i|S, \boldsymbol{p}(S))$ that a customer selects offer O_i from offer set S is thus:

$$Pr(O_i|S, \boldsymbol{p}(S)) = Pr(u_i > u_j)\ \forall\, j \in S \cup \{\emptyset\}; j \neq i \qquad (2)$$

If all offers in S have a negative utility, the customer chooses not to book by selecting the empty offer $\emptyset$ that has a utility of zero.

Optimization of offer set selection and pricing

The *expected net revenue contribution* of offer O_i in offer set S is computed by multiplying the offer's net revenue $(p_i - c_i)$ by the probability $Pr(O_i|S, \boldsymbol{p}(S))$ that a customer will select that offer from the offer set. The expected net revenue contribution of the entire offer set $R(S, \boldsymbol{p}(S))$ is the sum of the expected net revenues of all offers in the offer set:

$$R(S, \boldsymbol{p}(S)) = \sum_{O_i \in S} (p_i - c_i) * \Pr(O_i|S, \boldsymbol{p}(s)) \qquad (3)$$

The airline offer optimization problem is to select the optimal offer set S^* and correspondingly the optimal offer prices $\boldsymbol{p}^*(S^*)$ that maximize expected net revenue $R(S, \boldsymbol{p}(S))$ for a given search request. The optimal offer set and prices may depend on characteristics of the customer or the search request.

This problem can be separated into two distinct optimization problems:

- Offer Set Pricing: For a given offer set S, which offer prices $\boldsymbol{p}^*(S)$ maximize the offer set's expected net revenue $R(S, \boldsymbol{p}(S))$?
- Offer Set Selection: Which offer set $S^* \in \Omega$ maximizes the airline's expected net revenue from the search request?

The Offer Set Pricing problem can be represented mathematically as follows:

$$\boldsymbol{p}^*(S) = \underset{\boldsymbol{p}(S)}{\operatorname{argmax}}\, R(S, \boldsymbol{p}(S)) = \underset{\boldsymbol{p}(S)}{\operatorname{argmax}} \sum_{O_i \in S} (p_i - c_i) * \Pr(O_i|S, \boldsymbol{p}(S)) \qquad (4)$$

Given each offer set's optimal prices $\boldsymbol{p}^*(S)$ from the Offer Set Pricing problem, the Offer Set Selection problem identifies the offer set with the highest expected net revenue:

8

$$S^* = \underset{S \in \Omega}{\operatorname{argmax}} R\big(S, \boldsymbol{p}^*(S)\big) \tag{5}$$

Note that an airline could *effectively* perform Offer Set Selection by always offering the full offer set (consisting of all possible offers) and pricing the undesirable (to the airline) offers sufficiently high, such that customers would not purchase these unattractive offers. However, we believe that Offer Set Selection is an important second step: airlines face display space and bandwidth constraints that limit the offers that can be displayed, and may have a desire to provide customers with appealing content. In practice, Offer Set Selection is already performed today: by controlling availability of Basic Economy, airlines choose to make the brand unavailable (instead of pricing it unattractively) when they prefer to only sell their other economy brands.

In practice, airlines employ a variety of pure bundling, mixed bundling, and *a la carte* pricing for ancillary services. In recent years, there has been a trend of unbundling checked baggage from the flight itself. Many airlines charge an *a la carte* price for the bag, whereas others also offer a discounted price at the time of booking (mixed bundling) and a few still include it with all tickets (pure bundling). Yet there is no clear trend towards unbundling for other ancillary services. While carry-on bags are commonly offered *a la carte* at ultra low-cost carriers as of 2021 (i.e. Spirit Airlines and Ryanair), they remain bundled with the fare at most legacy airlines (i.e. Delta Air Lines and Air Canada). Some, like American Airlines, reverted their initial policy of prohibiting carry-on bags with basic economy fares and now bundle them with the fare.

Relationship between Dynamic Offer Generation, Dynamic Pricing, and traditional RMS

[Table 1 about here]

Table 1 summarizes how Dynamic Offer Generation differs from traditional RMS and dynamic pricing of the flight seat and ancillary services. With traditional RMS, prices are selected for the flight itinerary from a pre-defined set of filed fares. The price selected for an itinerary do not vary by customer segment. Ancillary services are also statically priced and do not vary by customer segment. Finally, flight itineraries and ancillary services are not bundled together; ancillary services are offered as optional purchases.

With dynamic pricing of the flight seat and ancillary services, the prices of the flight itinerary and ancillary services are selected from a continuous range of possible price points instead of from a set of pre-determined price points. These prices can also vary by customer segment. However, no bundling of flight itineraries and ancillary services is performed; in essence, the airline's offer set always allows the ancillary services to be purchased *a la carte*. We therefore refer to this strategy as Dynamic *A La Carte* pricing.

With Dynamic Offer Generation, flight itineraries and ancillary services may be bundled together into offers. The prices of the flight itinerary and ancillary services are optimized jointly when determining the offer price. These prices are determined as a function of the contents of each offer, as well as the other offers in the offer set. As we show below, the pricing and selection of the optimal offer set can also differ by customer segment.

Comparing DOG to traditional RMS and Dynamic *A La Carte* strategies allows us to investigate the benefits of the bundling component of DOG separately from the dynamic pricing of the flight and ancillary services.

Example: One flight itinerary, one ancillary service

Consider the example in Figure 3 with one flight itinerary f_1 and one ancillary service a_1. There are two possible offers, $O_1 = \{f_1\}$ and $O_2 = \{f_1, a_1\}$. Offer O_1 contains only the standalone flight, whereas offer O_2 contains both the flight and the ancillary service.

[Figure 3 about here]

We consider two possible offer sets, $S_1 = \{O_1, O_2\}$, and $S_2 = \{O_2\}$, with associated prices $\boldsymbol{p}(S_1) = (p_1, p_2)$ and $\boldsymbol{p}(S_2) = (p_2)$. In the "*a la carte* offer set" S_1, the customer can choose to purchase the flight alone (offer O_1) or pay more to purchase the offer with the ancillary service (offer O_2). In the "bundle offer set" S_2, the ancillary service is always bundled with the flight, and the customer cannot choose to purchase the flight alone. Note that, in general, the price p_2 for offer O_2 differs between offer sets S_1 and S_2. To resolve the notational conflict, we thus rewrite p_2 in offer set S_1 as $p_2 = p_1 + p_+$, where p_+ is effectively the additional fee for the ancillary service.

Figure 4 graphically illustrates how the customer's choice of offer depends on the offer set shown, the customer's WTPs θ_f and γ_a, and the offer prices p_1, p_+, and p_2. In the lightly shaded regions where WTPs are lower than the prices, the customer chooses not to book ($\emptyset$). In the gray regions, the customer purchases both the flight and the ancillary service because their total WTP $w_i = \theta_f + \gamma_a$ exceeds the price of the offer O_2. In the black regions of *a la carte* offer set S_1, customers do not purchase the ancillary service because $\gamma_a < p_+$, but they do purchase the flight since $\theta_f > p_1$. Note that a subset of these customers would have purchased the ancillary service if they were shown the bundle offer set S_2 instead.

[Figure 4 about here]

Figure 4 visualizes the probability $Pr\big(O_i \big| S, \boldsymbol{p}(S)\big)$ that a random utility-maximizing customer will purchase an offer. With independent WTPs θ_f and γ_a, these purchase probabilities can be written as:

$$S_1\!:\begin{cases} Pr\big(O_1\big|S_1,\boldsymbol{p}(S_1)\big) = Pr(\gamma_a < p_+) * \Pr(\theta_f > p_1) \\ Pr\big(O_2\big|S_1,\boldsymbol{p}(S_1)\big) = Pr(\gamma_a > p_+) * \Pr(\theta_f + \gamma_a > p_1 + p_+ | \gamma_a > p_+) \end{cases} \quad (6)$$
$$S_2\!:\big\{Pr(O_2|S_2,\boldsymbol{p}(S_2)) = Pr\big(\theta_f + \gamma_a > p_2\big)$$

The probabilities above can be interpreted as follows: if a customer is shown Offer Set S_1, he will choose to purchase Offer O_1 (flight alone) if he is willing to purchase the flight ($\theta_f > p_1$) but is does not value the ancillary enough to add it *a la carte* ($\gamma_a < p_+$).

10

He will purchase Offer O_2 (flight plus ancillary) if he desires the ancillary enough to purchase it *a la carte* ($\gamma_a > p_+$) and the sum of his valuation for both services exceeds the price for the flight and ancillary $\left(\theta_f + \gamma_a > p_1 + p_+ \middle| \gamma_a > p_+\right)$. If he is shown Offer Set S_2, he will purchase the bundled Offer O_2 if the sum of his valuation for both products exceeds the bundle price ($\theta_f + \gamma_a > p_2$).

Extension: One flight itinerary, two ancillary services

In our example with one flight and one ancillary service, we significantly reduced the problem size. As we later describe in *Curse of Dimensionality*, the number of offers scales rapidly with the number of ancillary services ($I = 2^K$), of which airlines often have a large catalog. The number of offer sets further escalates the problem size ($M = 2^I - 1$). The expressions for conditional purchase probability become complex for larger problems and are not easily solved without heuristic approximations.

For a catalog of one flight f_1 and two ancillary services a_1, a_2, there are four possible offers: $O_1 = \{f_1\}, O_2 = \{f_1, a_1\}, O_3 = \{f_1, a_2\}, O_4 = \{f_1, a_1, a_2\}$. As a result, there are $M = 15$ possible offer sets. Depending on business requirements, some offer sets may be more relevant than others. For example, an airline may be constrained by the number of offers that can be displayed in the distribution channel and perform Offer Set Selection among offer sets with cardinality k. Another conceivable constraint is that the comprehensive offer O_4 has to be included in all evaluated offer sets, as it may not make commercial sense to restrict passengers from buying all ancillary services.

The choice model naturally extends to larger offer sets. For example, for the largest offer set $S = \{O_1, O_2, O_3, O_4\}$, the utility-maximizing customer would purchase each offer with the following probabilities:

$$S: \begin{cases} Pr\left(O_1 \middle| S, \boldsymbol{p}(S)\right) = \Pr\begin{pmatrix} \left(\theta_f > p_1\right) \wedge \left(\gamma_{a1} < p_2 - p_1\right) \\ \wedge \left(\gamma_{a2} < p_3 - p_1\right) \wedge \left(\gamma_{a1} + \gamma_{a2} < p_4 - p_1\right) \end{pmatrix} \\ Pr\left(O_2 \middle| S, \boldsymbol{p}(S)\right) = \Pr\begin{pmatrix} \left(\theta_f + \gamma_{a1} > p_2\right) \wedge \left(\gamma_{a1} > p_2 - p_1\right) \\ \wedge \left(\gamma_{a2} < p_4 - p_2\right) \wedge \left(\gamma_{a1} - \gamma_{a2} > p_2 - p_3\right) \end{pmatrix} \\ Pr\left(O_3 \middle| S, \boldsymbol{p}(S)\right) = \Pr\begin{pmatrix} \left(\theta_f + \gamma_{a2} > p_3\right) \wedge \left(\gamma_{a1} > p_4 - p_3\right) \\ \wedge \left(\gamma_{a2} < p_3 - p_1\right) \wedge \left(\gamma_{a1} - \gamma_{a2} < p_2 - p_3\right) \end{pmatrix} \\ Pr\left(O_4 \middle| S, \boldsymbol{p}(S)\right) = \Pr\begin{pmatrix} \left(\theta_f + \gamma_{a1} + \gamma_{a2} > p_4\right) \wedge \left(\gamma_{a1} > p_4 - p_3\right) \\ \wedge \left(\gamma_{a2} > p_4 - p_2\right) \wedge \left(\gamma_{a1} + \gamma_{a2} > p_4 - p_1\right) \end{pmatrix} \end{cases} \quad (7)$$

The expressions above would simplify further with an assumption of independence between $\theta_f, \gamma_{a1}, \gamma_{a2}$ as well as an *a la carte* price constraint, namely $p_4 = p_3 + p_2 - p_1$, where there is no discount provided for purchasing both ancillary services together in offer 4.

Sensitivity analysis

Figure 5 shows the optimal offer prices p_1^* and p_2^* that result from the single-ancillary offer optimization problem. We assume a single customer type with flight WTP θ_f that is

11

Normally distributed with mean μ_f and standard deviation σ_f, and ancillary WTP γ_a that is Normally distributed with mean μ_a and standard deviation σ_a.

In the left panel of Figure 5, we vary the mean flight WTP μ_f from \$10 to \$400 and bid price π_f from \$0 to \$200, while holding the ancillary-related parameters constant at $m_a = \$20$, $\mu_a = \$25$, and $\sigma_a = \$7.5$. In the right panel, we vary the ancillary-related parameters from \$0 to \$70, while holding the flight-related parameters constant at $\pi_f = \$50$, $\mu_f = \$200$, and $\sigma_f = \$60$. In all cases, as μ_f and μ_a change, the standard deviations σ_f and σ_a scale as 30% of the mean.

[Figure 5 about here]

As seen in the left panel, the prices p_1^* and p_2^* increase with both the mean flight WTP μ_f and the bid price π_f. The increase is nearly linear where $\mu_f > \pi_f$. When $\mu_f < \pi_f$, the prices are generally bounded from below by their cost of provision ($p_1^* > \pi_f; p_2^* > \pi_f + m_a$). Given that passengers are expected to be willing to pay more for the bundle that includes the ancillary service, p_2^* is consistently more expensive than the flight price p_1^*. In the right panel, note that the flight price is nearly, but not completely, independent of the ancillary parameters μ_a and m_a, but the bundle price p_2^* shows a strong dependence on both sets of parameters.

Next, we compare the expected net revenue of the a la carte (S_1) and bundle (S_2) offer sets at the optimized prices $\boldsymbol{p}^*$ to understand when bundle pricing is expected to generate higher revenue for the airline. In Figure 6, this condition is indicated by the bundle surface laying above the a la carte surface. We observe that the flight-related parameters have a very limited effect on the offer set selection, as the expected net revenue is very similar for the two offer sets. Nonetheless, we observe that when the bid price π_f is low, the algorithm has a slight preference for the a la carte offer set. In the remaining regions, the bundle offer set is preferred.

[Figure 6 about here]

On the other hand, ancillary-related parameters strongly influence the offer set selection (Figure 6). When μ_a is sufficiently higher than the cost m_a, bundling is the preferred pricing strategy, providing higher expected net revenue than the a la carte offer set by a small margin. But when $\mu_a < m_a$, a la carte pricing significantly outperforms bundling. Figure 6 shows a decision boundary where both offer sets have the same $R(S, \boldsymbol{p}^*(S))$. This boundary is relatively consistent across a variety of tested μ_f and π_f. It can be linearly approximated by the following rule of thumb:

$$\text{If } \gamma_a \sim \mathcal{N}(\mu_a, (0.3 * \mu_a)^2), \text{ then:}$$
$$R(S_2, \boldsymbol{p}^*(S_2)) > R(S_1, \boldsymbol{p}^*(S_1)) \Rightarrow \mu_a \geq 1.25 * m_a$$

12

These observations have implications for dynamic offer generation: For a segment of passengers whose ancillary WTP can be modeled with a normal distribution, the selection of the optimal offer set (i.e. bundle vs *a la carte*) is often independent of the flight price. This means that for any one segment, the offer set selection might not need to be very dynamic during the selling horizon and bundling could be realizable through simple business rules.

13

PERFORMANCE OF DOG IN THE PASSENGER ORIGIN-DESTINATION SIMULATOR (PODS)

Simulation setup

Dynamic offer generation was implemented and tested in the Passenger Origin-Destination Simulator (PODS), a multi-agent simulation of the passenger choice process in a competitive airline network. For a more detailed description of the PODS simulator and its models, see Bockelie (2019). All tests were conducted in a hypothetical network with four competing airlines, each operating a primarily hub-and-spoke network with a total of 442 flights per day serving 572 origin-destination markets. All airlines file a fully unrestricted fare structure in each market, like fare structures commonly found within Europe today, with no advance purchase requirements.[1] All airlines use sell-up WTP forecasting (Belobaba and Hopperstad, 2004) and marginal revenue fare adjustment (Fiig et al., 2010) with a network RM optimizer based on Displacement Adjusted Virtual Nesting (Smith and Penn, 1988).

We test DOG with one ancillary service that is modeled after a checked bag. As in Figure 3, DOG can decide for every search request whether to offer the ancillary *a la carte* for an additional fee (offer set S_1), or whether to bundle it with the flight at a single price (offer set S_2). We provide a summary of the implementation of DOG in PODS below; readers interested in a more detailed description can refer to Wang (2020).

In our simulations, leisure passengers are willing to pay more than business passengers for the ancillary service, whereas business passengers have higher mean flight WTP μ_f. Following the "Q-multiplier" heuristic proposed by Wittman and Belobaba (2018), we approximate flight WTP for a passenger segment as Normally distributed with mean $\mu_f = Qmult * P_L(m)$ and standard deviation $\sigma_f = \alpha * \mu_f$, where $P_L(m)$ is the lowest filed fare by the airline in market m.

Half of all business passengers draw their ancillary WTP from a Normal distribution with $\mu_a = \$25$ and $\alpha = 0.3$, while the other half does not require checked baggage ($\gamma_a = \$0$). All leisure passengers have a Normally distributed ancillary WTP with $\mu_a = \$31$ ($\alpha = 0.3$), which is roughly 25% higher than that of business passengers. All airlines incur a cost of $m_a = \$25$ per ancillary sold.

We assume that all DOG airlines can correctly classify a search request to its customer segment. They also accurately estimate both passenger types' mean ancillary WTP μ_a. They approximate the passengers' flight WTP μ_f using constant Q-multipliers of 2.7 and 1.1 for business and leisure segments respectively. The estimation of the WTP parameters from historical data is outside the scope of this paper; maximum likelihood estimation

[1] Unrestricted fare structures are also common in North America, though often with advance purchase requirements. Fare adjustment, part of some RM systems, can replicate the function of advance purchase requirements.

14

methods similar to those proposed in Fiig et al. (2014) and Newman et al. (2014) could be used to estimate these parameters.

In our simulations, we bound the offer prices from DOG to the L filed fares from the airline's traditional pricing structure $\boldsymbol{P} = (P_1, \dots, P_L)$. This allows the airline to maintain the hierarchy of booking classes in the RMS. Bounding DOG prices to traditional prices could also be desirable in a competitive market, where any large price disturbances might lead to competitive responses by other airlines.

Let us define P_c as the filed fare of the lowest available fare class in a traditional RMS, P_{c-1} as the filed fare of the next higher (more expensive) fare class, and P_{c+1} is the filed fare of the next lower (unavailable) fare class. If $P_c = P_L$ (the lowest-priced fare class), we set $P_{c+1} = P_c$ and if $P_c = P_1$ (the highest fare class), then $P_{c-1} = P_c$.

In our modeling, the flight price in offer set S_1, $\tilde{p}_1$, is bounded to be within $\frac{P_c+P_{c+1}}{2} \leq \tilde{p}_1 \leq \frac{P_c+P_{c-1}}{2}$, such that every price point $\tilde{p}_1$ is uniquely mapped to a single fare class in the traditional pricing structure. The prices of all other offers are bounded based on $\tilde{p}_1$ to maintain the relative price difference $p_i^* - p_1^*$ of the original unbounded prices p_1^*:

$$\tilde{p}_i = \tilde{p}_1 + (p_i^* - p_1^*) \;\forall\; 1 < i \leq I \tag{8}$$

We illustrate the bounding with a numerical example in Figure 7: Suppose the optimized prices from the DOG price optimization are $p_1^* = \$170$ and $p_2^* = \$193$ for offer sets S_1 and S_2, respectively. If $P_{c-1} = \$260$, $P_c = \$200$ and $P_{c+1} = \$160$, the flight price will be bounded to $\tilde{p}_1 = \$180$. Then the bundle price will be $\tilde{p}_2 = \$180 + (\$193 - \$170) = \$203$. The selection of the optimal offer set S^* is then performed with the bounded prices $\tilde{\boldsymbol{p}}(S)$.

[Figure 7 about here]

We compare DOG to the traditional RMS and Dynamic *A La Carte* pricing strategies. Recall from Table 1 that an airline using either strategy offers only the *a la carte* offer set $S_1 = \{O_1, O_2\}$, and sets the price of the ancillary service p_+ independently of the flight itinerary price p_1. With traditional RMS, p_1 is selected among the pre-filed fares $\boldsymbol{P}$ and p_+ is fixed at \$33.59, which myopically maximizes ancillary net revenue under the 39%/61% mix of business and leisure bookings observed in the simulation. In Dynamic *A La Carte* pricing, p_1 is segmented by passenger type and dynamically priced using the DOG heuristic. p_+ is also segmented and priced to maximize ancillary net revenue at \$30.64 or \$34.10 for business and leisure passengers, respectively.

Simulation Results

Table 2 summarizes the impact of dynamic offer generation on one airline's (Airline 1) net revenue (flight and ancillary revenue net of ancillary cost of provision), net yield (net revenue divided by revenue passenger miles) and load factor (revenue passenger miles divided by available seat miles). When only Airline 1 uses DOG, its net revenue increases by +4.6% and its load factor increases by +0.6pts. If all four airlines switch from traditional RMS to DOG, Airline 1's net revenue gain reduces to +1.7% with a load factor increase of +0.2pts. DOG is therefore revenue positive in both asymmetric and symmetric tests.

[Table 2 about here]

Figure 8 shows the impact of DOG on the total net revenue of all four airlines in the simulation. When only Airline 1 uses DOG, its gain comes at the expense of the remaining airlines using traditional RMS. When all airlines use DOG, the competitive benefit for Airline 1 is reduced and all airlines increase their revenue over the traditional baseline.

[Figure 8 about here]

We also show equivalent results for Dynamic *A La Carte* pricing, where no bundles are offered but the flight seat and ancillary services are dynamically priced. Dynamic *A La Carte* results in a smaller +3.7% revenue gain for Airline 1 when it is the first mover, but a +1.8% revenue gain (comparable to DOG) when all airlines implement it. These results are similar to previously reported simulation studies on the value of segmented dynamic pricing of the flight seat (Wittman and Belobaba, 2018). The results also show that the bundling of the ancillary service drives an incremental +0.9% revenue benefit, when it is used as a competitive tool to offer a lower total price than competitors selling ancillaries *a la carte*.

[Figure 9 about here]

Figure 9 shows the breakdown of offers purchased by passengers of Airline 1. With traditional RMS (No DOG), 26% of passengers choose to buy the ancillary service (representing 39.0% of leisure passengers and 6.3% of business passengers). When all airlines use DOG (All DOG), the ancillary service was already bundled in the offer selected by 36% of passengers. The remaining 63% of passengers could purchase the service at a segmented price p_+^* of $30.64 for business passengers or $34.10 for leisure passengers.

With our choice of ancillary WTP parameters, all business passengers are offered the a la carte offer set S_1, which maximizes expected net revenue given the low WTP for the ancillary service. Meanwhile, 60% of leisure passengers were offered the bundle, which generally occurred at higher flight prices. Overall, 41% of AL1's passengers purchased the bundle when it was the only airline to offer bundles, but only 36% of passengers do so once other airlines offered bundles as well. This suggests that the primary benefit of dynamic bundling is as a competitive tool against *a la carte* competitors.

These values appear reasonable: Shao and Kauermann (2020) report observed purchase rates for a European airline's branded fares that are primarily distinguished through their

16

checked baggage allowances. In a mixed bundling scenario where passengers could either purchase the ancillary service bundled with the flight *or* separately from the flight, 47-58% of passengers chose to purchase the bundled fare in their study, while an additional 9-15% of passengers purchased the *a la carte* fare and added the ancillary service later.

[Figure 10 about here]

Figure 10 shows the relative change in average net revenue (price paid less ancillary costs) from AL1's business and leisure passengers as it switches from traditional RMS to DOG. Note that the flight prices generally increase towards the day of departure, which is not visible in the relative change. DOG prices differ from those of traditional RMS due to the price segmentation across business and leisure passengers. With our choice of Q-multipliers used in DOG, leisure flight prices decrease and business prices generally increase compared to traditional RMS.

[Figure 11 about here]

As the segmented pricing of the flight (Dynamic *A La Carte*) contributes the majority of the revenue benefit of DOG, the Q-multipliers as estimates of each passenger segment's flight WTP θ are the most important parameters of the model. As shown in Figures 11 and 12, the revenue gain of DOG is sensitive to both the business and leisure Q-multipliers used in the simulation as they directly influence the flight price offered to the customers. In general, we observe that increasing the Q-multiplier decreases Airline 1's revenue in asymmetric scenarios (its flights become more expensive compared to the unchanged competitors). On the other hand, the revenue increases for all airlines when all competitors match the higher Q-multipliers in the symmetric tests.

[Figure 12 about here]

Figure 12 shows clearly that AL1's revenue is extremely sensitive to the leisure Q-multiplier in competitive, asymmetric scenarios. In the simulation, large market share shifts can occur when Airline 1's leisure pricing significantly differs from that of its three competitors. In the symmetric tests, we note that revenue does not increase further when prices are increased beyond a Q-multiplier of 1.2, because of the limited flight WTP of leisure passengers in PODS.

[Figure 13 about here]

DOG assumes that all passengers act rationally by maximizing their utility across flights and ancillary services. In Figure 13, we study the performance of DOG when passengers in PODS are instead simulated under the *sequential* choice model of Bockelie and Belobaba (2017). Sequential passengers first choose their flight itinerary without considering any ancillary services, and then evaluate the purchase of ancillary services independently at a second stage (i.e. when checking in for the flight). Therefore, they do not consider the presence of included ancillary services when evaluating offers across airlines and will choose the lowest flight price, even if the airline charges for the ancillary service and it results in a higher total price compared to a competing bundled offer.

17

We observe that the revenue benefit of DOG diminishes as more passengers evaluate ancillary services sequentially. When only Airline 1 uses DOG, it underperforms the traditional RMS baseline by -1.9% when all simulated passengers are sequential. On the other hand, with Dynamic *A La Carte* pricing, Airline 1 revenue increases by +3.7% independent of the proportion of sequential passengers in the simulation. Even when selling the ancillary *a la carte*, Airline 1 benefits from segmented, dynamic flight pricing that undercuts the traditional prices of competing airlines. This benefit of DOG's price optimization persists independent of ancillary choice behavior. But DOG underperforms overall when it shows the bundled offer set to sequential passengers as it overestimates the sequential passenger's valuation for bundles compared to the a la carte offers of the competitor airlines.

DISCUSSION AND MANAGERIAL IMPLICATIONS

Our heuristics show how prices of offers that combine flights with ancillary services can be optimized in an offer set, with intuitive results: the price of an offer generally increases with the cost of providing it, as well as the passenger's willingness-to-pay for it, and the optimal price for one offer depends on the prices of other offers within the offer set. Despite this relative simplicity and intuitiveness, DOG is a fundamental change for revenue management. The technological difficulties and human challenges of integrating and successfully managing a new dynamic offer generation engine should not be underestimated. We list five key considerations here:

1. *Competition*: Our simulations did not include competitive responses to the new dynamic pricing structure and bundling. In reality, traditional airlines could adjust their fare structures, bundling strategies, or fare class availability decisions to protect their own market share. Depending on the type and magnitude of competitive response, the revenue gain realized with DOG could be reduced.

2. *Bundling Strategy*: The dynamic bundling in DOG introduces an explicit trade-off between pursuing more *a la carte* offers, to increase ancillary revenue, or more bundled offers, to increase flight revenue. A comprehensive view of revenue streams, and an understanding of the new mechanisms to control bundling, is necessary to select the strategy that increases *total* revenue. To reduce the initial change management, an implementation could start with Dynamic *A La Carte* pricing, capturing a significant portion of the revenue benefit, before implementing full DOG. Reducing the initial implementation scope could also decrease time to market, providing a first mover advantage. In all of our tests, the revenue gains were higher for the DOG or Dynamic *A La Carte* airline when its competitors were still using traditional pricing and RM.

3. *Curse of Dimensionality*: With I possible offers, the number of offer sets is $2^I - 1$ if no restrictions are placed on which combinations of offers are permitted in an offer set. Furthermore, for an airline with one flight and K ancillary services, the number of possible offers itself also grows as 2^K, such that the number of possible

18

offer sets could be as high as $2^{2^K} - 1$. This would make the brute-force evaluation of $R(S)$ for all S at large K infeasible. To manage the size of the solution space, airlines could limit the number of ancillaries included in the DOG Offer Set Generation process to only the highest revenue products, while selling the remaining low-volume ancillaries through traditional *a la carte* channels. Alternatively, airlines could require that all offer sets consist of exactly n offers, corresponding to the number of offers that can be displayed in the distribution channel.

4. *Retailing Becomes a Priority*: Our simulations showed that bundling can increase revenue with utility-maximizing choice behavior, but can decrease it with sequential behavior. Customers must be aware that they are purchasing a bundle and must understand what the bundle contains for a bundling strategy to increase revenue – information that is traditionally difficult for the airline to communicate to customers. New technologies such as NDC, ATPCO's Amenities, Universal Product & Ticket Attributes, and Next Generation Storefront all allow airlines to publish rich content and descriptions of their offers to sales channels around the world; airlines will need to ensure that their websites and distribution partners are effectively leveraging this information. Customers who do not understand that an airline is selling a bundle will exhibit sequential behavior and may turn to a competitor's *a la carte* offer with a lower base fare, even if the final price, including desired ancillaries, is higher. In distribution channels where such transparency about offer contents and ancillary fees is not achievable, unbundled fares and a la carte offer sets yield better results.

5. *Breaking Silos*: Effective management of DOG will require close cooperation between revenue management and work groups who may not have interacted on a regular basis in the past, such as e-commerce, loyalty, and customer relationship management (CRM). As one example, customer segmentation models are often built and maintained by a CRM department, and flight and ancillary WTP parameters are controlled by RM & ancillary management departments. DOG requires an explicit integration of these areas as its input WTP parameters θ_j & γ_k are segment-specific. CRM will need to ensure its models meet the needs of revenue managers, and revenue managers will need to understand the customer segmentation strategy and embed it each step of their daily analysis workflow.

19

CONCLUSIONS

The transformation of the airline distribution process with the New Distribution Capability (NDC) and the growth of airline ancillary revenue streams are two major trends in the airline industry. In this paper, we presented a new optimization approach to generate customized offers and bundles. Together with a traditional RM system, dynamic offer generation (DOG) calculates segmented and continuous prices for both flights and ancillaries in pursuit of total revenue optimization. We used DOG to study the potential benefits of airline offer optimization and its implications for the airline industry.

Our simulations led to three important conclusions. First, the majority of revenue gains reported from DOG are attributable to the competitive effects of segmented, dynamic flight pricing and are also observed when no bundles are offered. Second, bundling can be a competitive advantage and further increase revenue by +0.9% in asymmetric tests with 100% utility-maximizing passengers. Finally, bundling is shown to be ineffective when passengers do not consider the price and value of ancillary services when choosing the airline and flight itinerary. Our insights provide valuable guidance to practitioners that are navigating the organizational and technological challenges of implementing offer optimization, which we explored in the previous section.

Our problem formulation is very general and can be extended to include more ancillary services or inventory constraints. Yet our implementation makes several assumptions. The heuristic assumes that customers choose offers that maximize their utility and the benefit of bundling diminishes when the distribution channel encourages them to behave sequentially instead. We also assume that airlines can estimate flight and ancillary willingness-to-pay from their historical data and accurately segment trip requests, which are all challenging problems that were outside the scope of this paper.

DOG is only a first step towards a world of offer optimization. In particular, the assortment optimization adds an entirely new dimension to the problem space of revenue management. We have shown that optimizing the offer set shown to customers can be a competitive advantage and lead to revenue gains. As such, the future of airline RM lies in selling the right *set of offers* to the right consumer at the right time and at the right prices.

ACKNOWLEDGEMENTS

The opinions expressed in this paper represent those of the authors only and do not necessarily reflect the positions of their respective employers. The authors are grateful to Peter Belobaba and the members of the MIT PODS Consortium for financial support and for their ideas that culminated in this paper, and to Matthew Berge for assisting with development. We also thank Thomas Fiig and Richard Cléaz-Savoyen for reviewing the paper and offering their thoughtful feedback.

20

REFERENCES

Adams, W.J. and J.L. Yellen. 1976. Commodity bundling and the burden of monopoly. *The Quarterly Journal of Economics* 90(3): 475-498.
Belobaba, P. and C. Hopperstad. 2004. Algorithms for revenue management in unrestricted fare markets. Presentation at *AGIFORS Reservations and Yield Management Study Group*.
Bockelie, A. 2019. "Ancillary services in the airline industry: Passenger choice and revenue management optimization." PhD thesis, Massachusetts Institute of Technology, https://hdl.handle.net/1721.1/121879.
Bockelie, A. and P. Belobaba. 2017. Incorporating ancillary services in airline passenger choice models. *Journal of Revenue and Pricing Management* 16: 553-568.
Cataldo, A. and J.-C. Ferrer. 2017. Optimal pricing and composition of bundles: A two-step approach. *European Journal of Operations Research* 259: 766-777.
Dezelak, M. and R. Ratliff. 2018. Towards new industry-standard specifications for air dynamic pricing engines. *Journal of Revenue and Pricing Management* 17: 394-402.
Dobson, G. and S. Kalish. 1988. Positioning and pricing a product line. *Marketing Science* 7.2: 107-125.
Ellison, G. 2005. A model of add-on pricing. *The Quarterly Journal of Economics* 120(2): 585-637.
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): 152-170.
Fiig, T., R. Härdling, S. Pölt, and C. Hopperstad. 2014. Demand forecasting and measuring forecast accuracy in general fare structures. *Journal of Revenue and Pricing Management* 13: 413-439.
Fiig, T., O. Goyons, R. Adelving, and B. Smith. 2016. Dynamic pricing – The next revolution in RM? *Journal of Revenue and Pricing Management* 15: 360-379.
Fiig, T., R. Le Guen, and M. Gauchet. 2018. Dynamic pricing of airline offers. *Journal of Revenue and Pricing Management* 17: 381-393.
Gabaix, X. and D. Laibson. 2006. Shrouded attributes, consumer myopia, and information suppression in competitive markets. *The Quarterly Journal of Economics* 121(2): 505-540.
Hanson, W. and R. K. Martin. 1990. Optimal bundle pricing. *Management Science* 36(2): 155-174.
Hao, E. 2014. "Ancillary revenues in the airline industry: impacts of revenue management and distribution." Master's thesis, Massachusetts Institute of Technology, http://hdl.handle.net/1721.1/89854.
Jagabathula, S. and P. Rusmevichientong. 2017. A nonparametric joint pricing and assortment optimization model. *Management Science* 63(9): 3128-3145.
Kumar, R., A. Li, and W. Wang. 2018. Learning and optimizing through dynamic pricing. *Journal of Revenue and Pricing Management* 17: 63-77.
Madireddy, M., R. Sundararajan, G. Doreswamy, M.H. Nia, and A. Mital. 2017. Constructing bundled offers for airline customers. *Journal of Revenue and Pricing Management* 16: 532 – 552.
Newman, J., M. Ferguson, L. Garrow, and T. Jacobs. 2014. Estimation of choice-based models using sales data from a single firm. *Manufacturing & Service Operations Management* 16(2): 184-197.
Ødegaard, F. and J.G. Wilson. 2016. Dynamic pricing of primary products and ancillary services. *European Journal of Operations Research* 251: 586-599.
Schmalensee, R. 1984. Gaussian demand and commodity bundling. *The Journal of Business* 57(1-2): 211-230.
Shao, S. and G. Kauermann. 2020. Understanding price elasticity for airline ancillary services. *Journal of Revenue and Pricing Management* 19: 74-82.
Shukla, N., A. Kolbeinsson, K. Otwell, L. Marla, and K. Yellepeddi. 2019. *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*: 2174-2182.

21

Shulman, J.D. and X. Geng. 2013. Add-on pricing by asymmetric firms. *Management Science* 59(4): 899-917.
Smith, B. C. and C. W. Penn. 1988. Analysis of alternate origin-destination control strategies. *Proceedings of the Twenty Eigth Annual AGIFORS Symposium*.
Talluri, K. and G. van Ryzin. 2004. Revenue management under a general discrete choice model of consumer behavior. *Management Science* 50(1): 15-33.
Vinod, B., R. Ratliff, and V. Jayaram. 2018. An approach to offer management: Maximizing sales with fare products and ancillaries. *Journal of Revenue and Pricing Management* 17: 91-101.
Wang, K.K. 2020. "Airline revenue management with dynamic offers: Bundling flights and ancillary services." Master's thesis, Massachusetts Institute of Technology, https://hdl.handle.net/1721.1/127110.
Westermann, D. 2013. The potential impact of IATA's New Distribution Capability (NDC) on revenue management and pricing. *Journal of Revenue and Pricing Management* 12(6): 565-568.
Wittman, M.D. and P. Belobaba. 2018. Customized dynamic pricing of airline fare products. *Journal of Revenue and Pricing Management* 17(2): 78-90.
Wittman, M.D. and P.P. Belobaba. 2019. Dynamic pricing mechanisms for the airline industry: a definitional framework. *Journal of Revenue and Pricing Management* 18: 100-106.

22

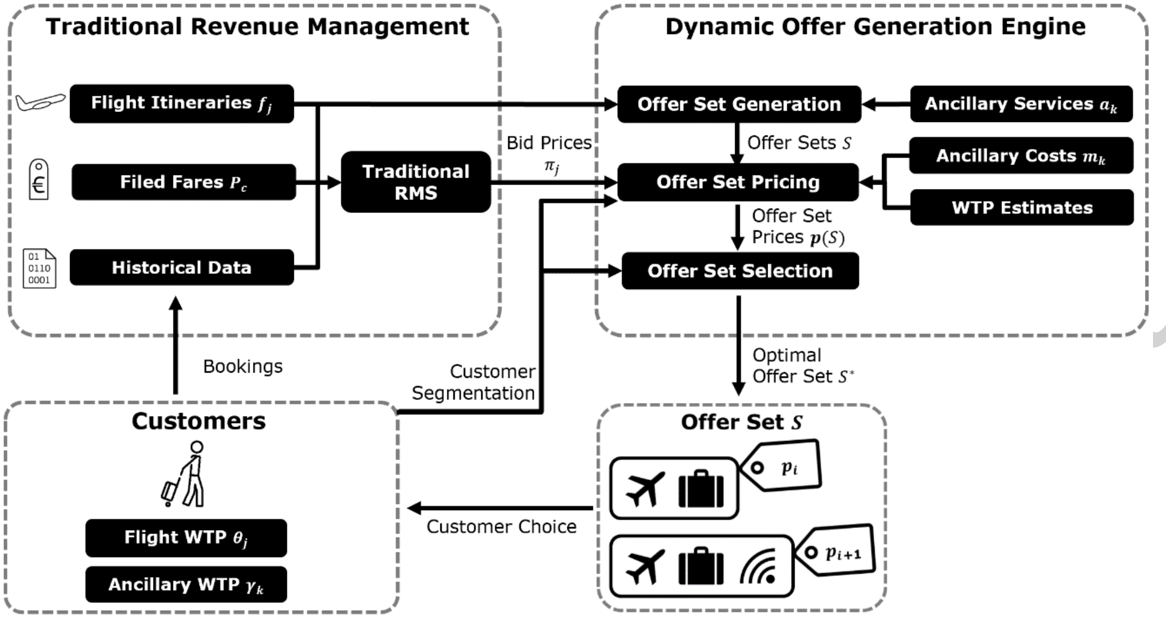

Figure 1: Schematic of interactions between Dynamic Offer Generation and a traditional revenue management system

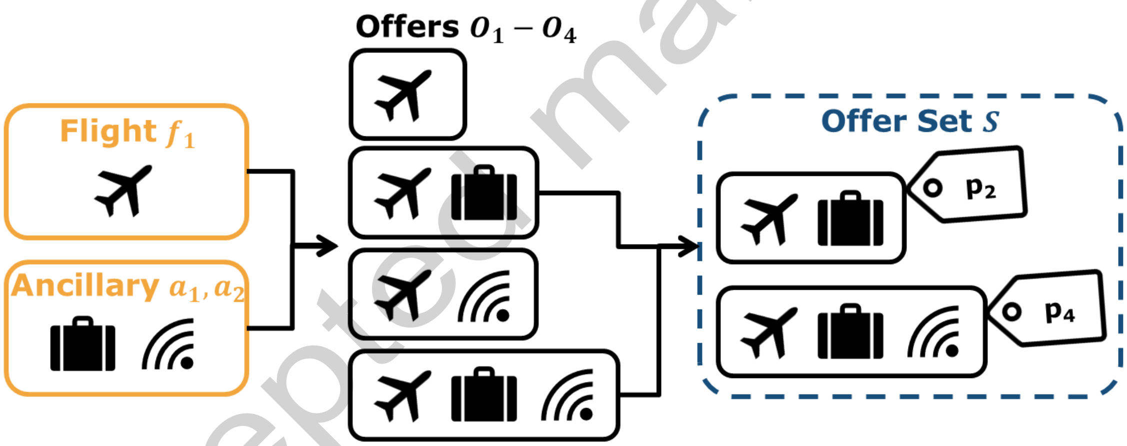

Figure 2: Illustration of an example offer set consisting of two offers $S = \{O_2, O_4\}$. In this bundled offer set, ancillary a_1 (baggage) is bundled with the flight f_1, while ancillary a_2 (internet) is an optional add-on.

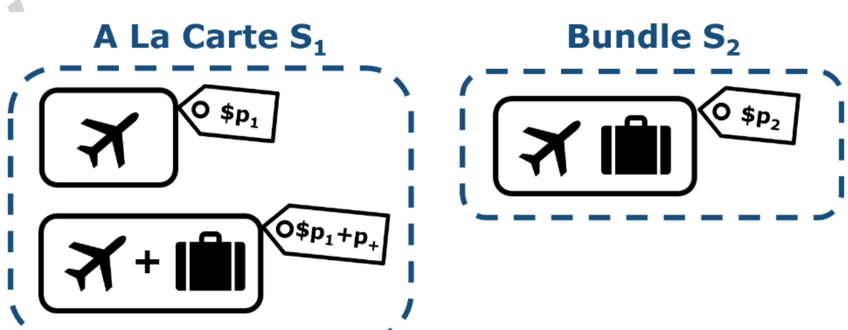

Figure 3: Composition and prices of offer sets S_1 and S_2 with one flight and one ancillary service.

23

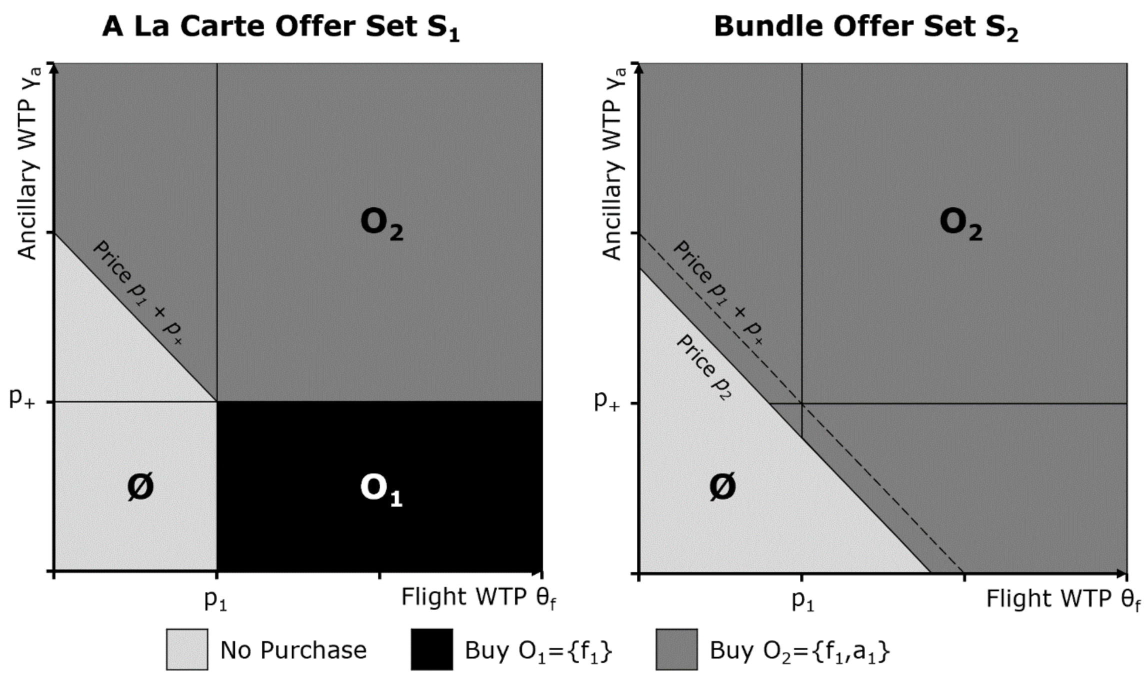

Figure 4: Selected offer from Offer Set S_1 (left panel) and Offer Set S_2 (right panel) as a function of customer flight WTP and ancillary WTP

24

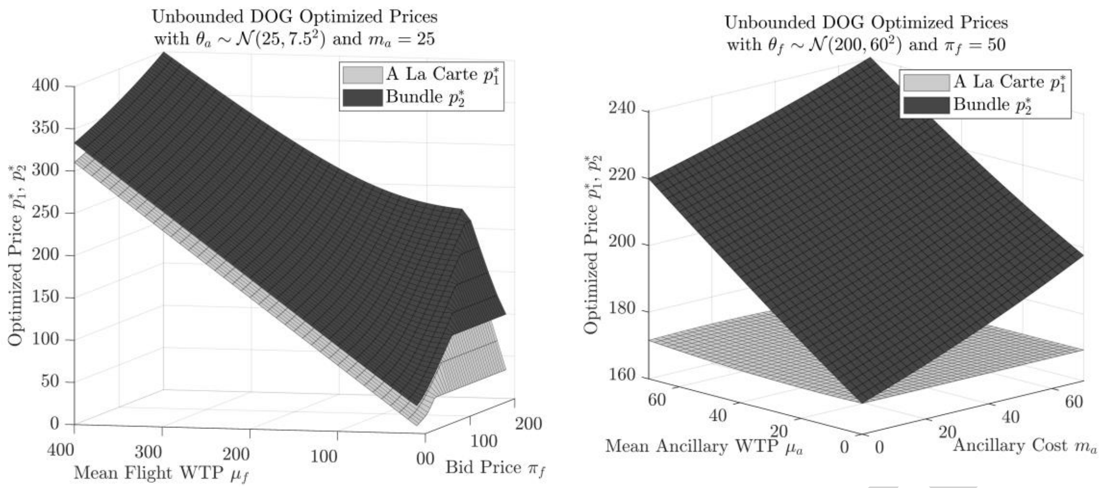

Figure 5: Optimal offer prices p_1^* and p_2^* in the single-flight, single-ancillary case as a function of mean flight WTP μ_f and bid price π_f (left panel) and mean ancillary WTP μ_a and ancillary cost m_a (right panel)

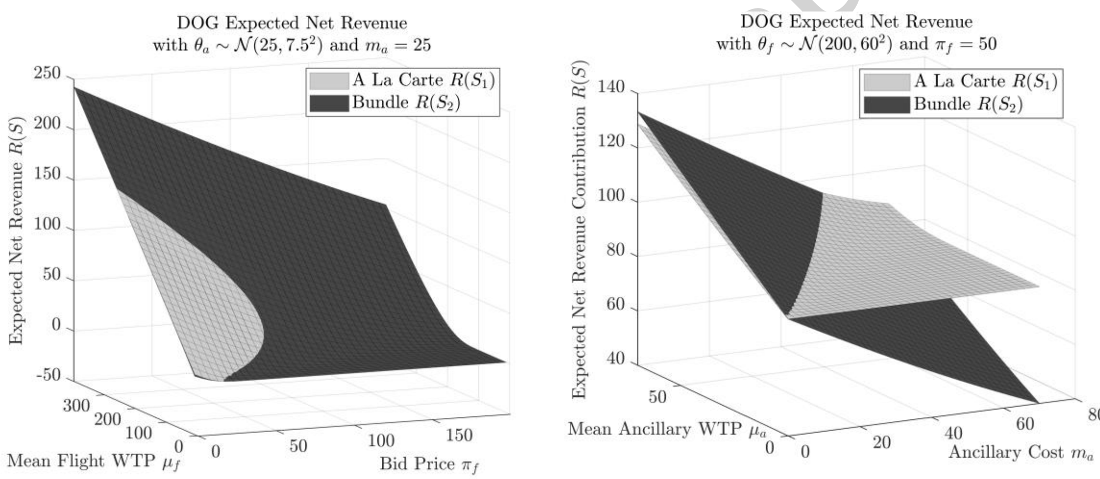

Figure 6: Expected net revenue of the a la carte and bundled offer sets in the single-flight, single-ancillary case as a function of mean flight WTP μ_f and bid price π_f (left panel) and mean ancillary WTP μ_a and ancillary cost m_a (right panel)

25

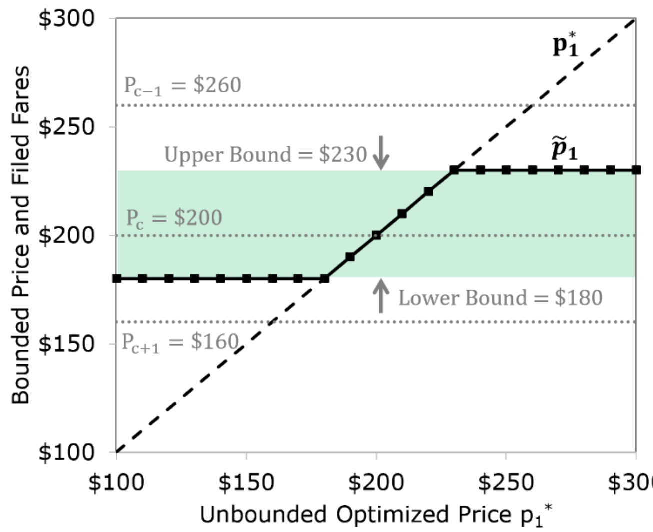

Figure 7: Illustration of the unbounded price p_1^* and the resulting bounded price $\tilde{p}_1$ within the shaded bounded price range.

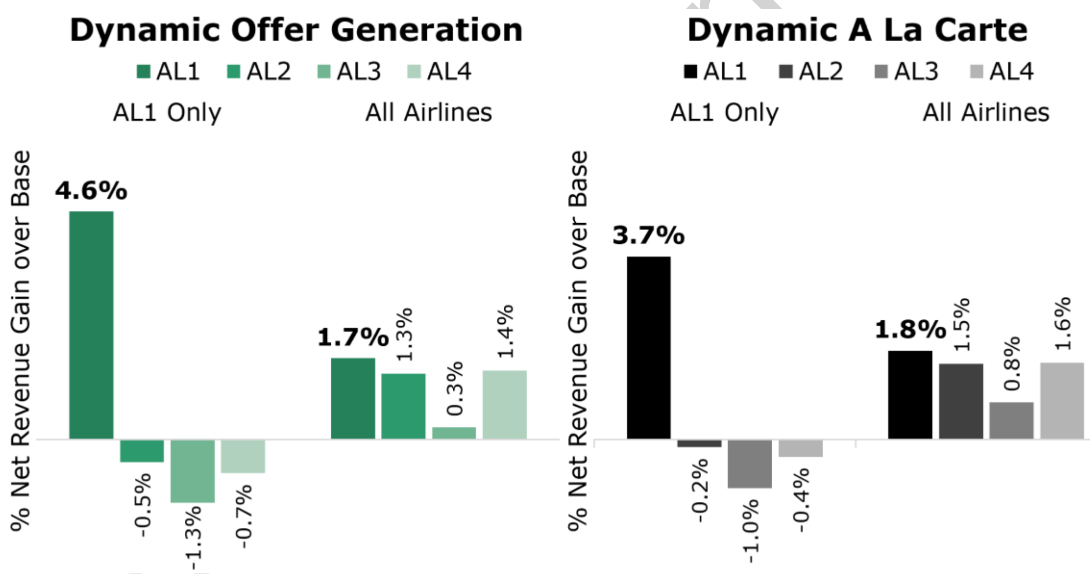

Figure 8: Change in net revenue when either only Airline 1 or all airlines shift from Traditional RMS to DOG (left panel) or Dynamic *A La Carte* (right panel)

26

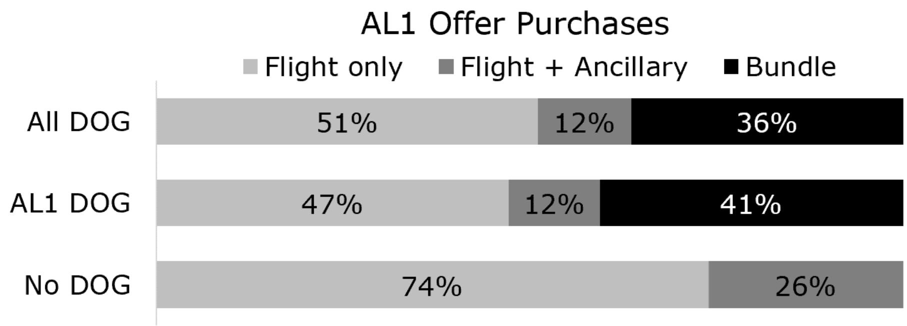

Figure 9: Offer purchases for Airline 1 when zero, one, or all airlines shift from Traditional RMS to DOG

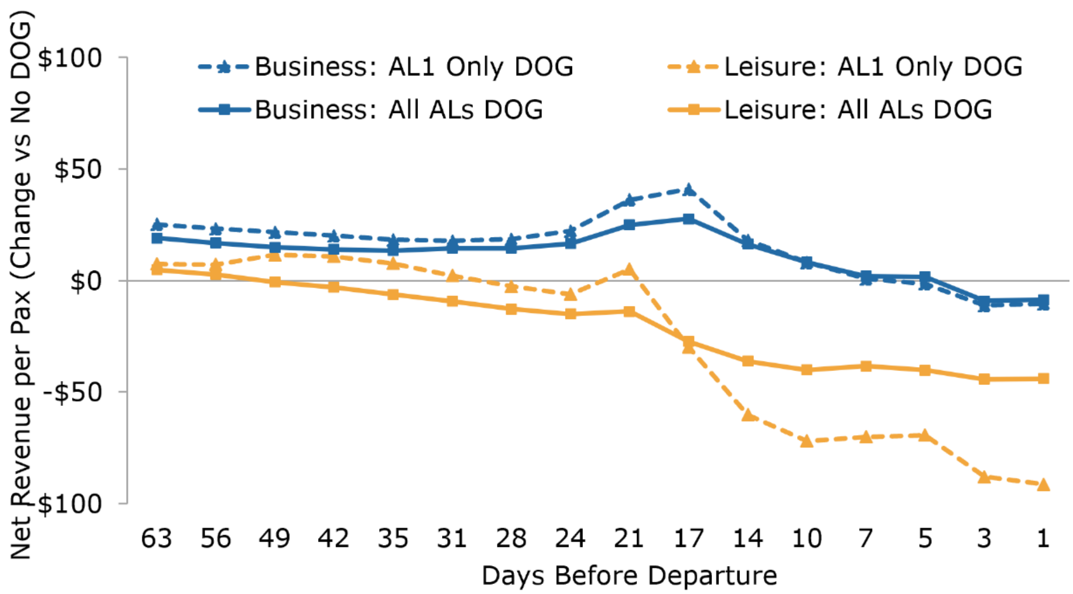

Figure 10: Change in Airline 1's average net revenue per passenger (business or leisure) by days before departure when one or all airlines shift from Traditional RMS to DOG

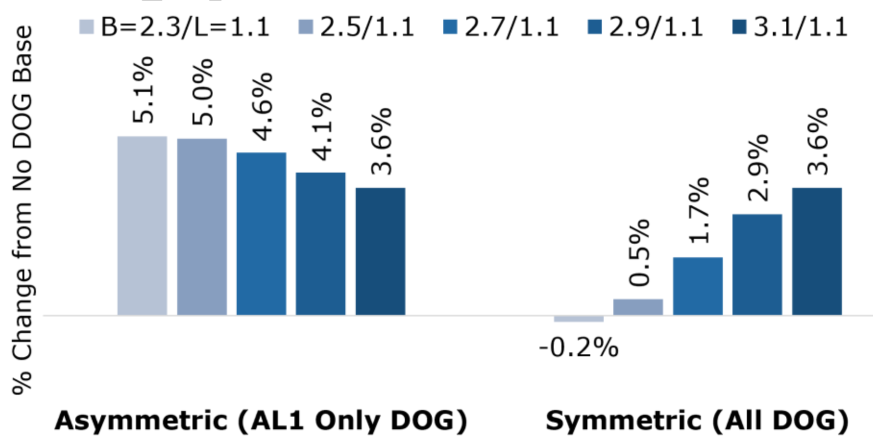

Figure 11: Airline 1's change in net revenue from traditional RM under different business Q-multipliers (θ estimates) and constant leisure Q multiplier of 1.1 (all DOG airlines use the same Q-multipliers)

27

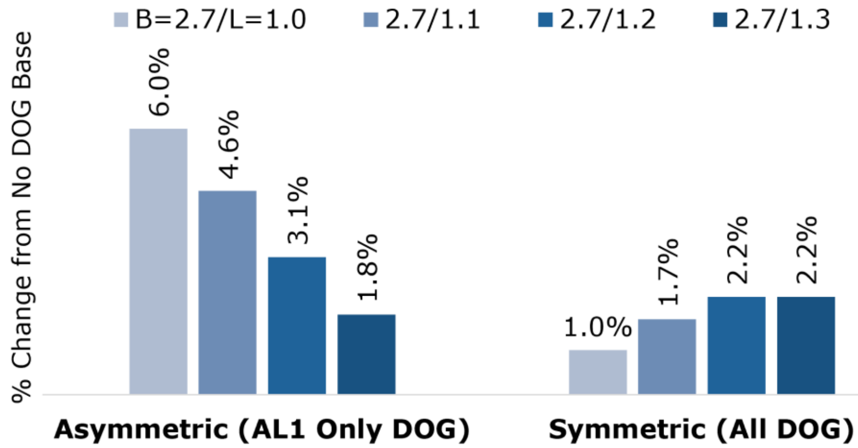

Figure 12: Airline 1's change in net revenue from traditional RM under different leisure Q-multipliers (θ estimates) and constant business Q multiplier of 2.7 (all DOG airlines use the same Q-multipliers)

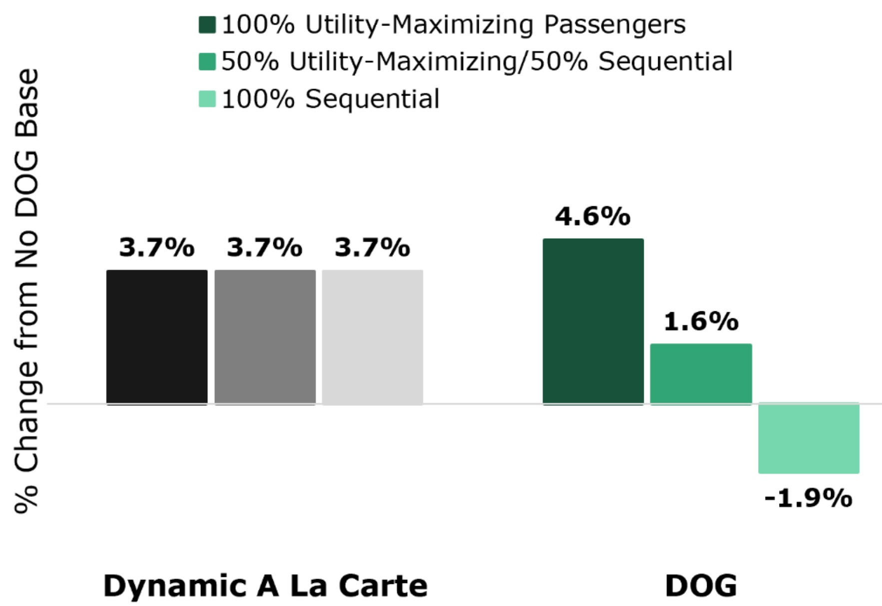

Figure 13: Airline 1's change in net revenue from traditional RM when only Airline 1 shifts from traditional RM to either Dynamic *A La Carte* pricing or DOG under different simulated ancillary choice behaviors

Strategy	Price Optimization		Bundled Offers?
	Flight Itinerary	Ancillary Services	
Traditional RMS	Filed fares, not segmented	Static, not segmented	No
Dynamic *A La Carte*	Continuous, segmented	Continuous, segmented	No
Dynamic Offer Generation	Continuous, segmented	Continuous, segmented	Yes

Table 1: Comparison of Dynamic Offer Generation to Traditional RMS and Dynamic *A La Carte* pricing

Airline 1	No DOG	AL1 Only DOG	All ALs DOG
Net Revenue	$2,858,341	$2,990,873 +4.6%	$2,905,600 +1.7%
Net Yield	$0.1218	$0.1266 +3.9%	$0.1235 +1.4%
Load Factor	82.06%	82.63% +0.6 pts	82.29% +0.2 pts

Table 2: Net revenue, net yield and load factor when one or all airlines shift from Traditional RMS to DOG

29