

MIT Open Access Articles

Dynamic offer creation for airline ancillaries using a Markov chain choice model

The MIT Faculty has made this article openly available. ***Please share***
how this access benefits you. Your story matters.

Citation: Wang, Kevin. 2023. "Dynamic offer creation for airline ancillaries using a Markov chain choice model." *Journal of Revenue and Pricing Management*.

As Published: 10.1057/s41272-022-00398-3

Persistent URL: <https://hdl.handle.net/1721.1/147886>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Creative Commons Attribution 4.0 International license





Dynamic offer creation for airline ancillaries using a Markov chain choice model

Kevin K. Wang¹ · Michael D. Wittman² · Thomas Fiig²

Received: 25 June 2022 / Accepted: 15 September 2022
© The Author(s) 2023

Abstract

Customers have become accustomed to a highly streamlined and personalized experience when shopping online. While tech giants such as Apple, Amazon, and Netflix are experts in using customer information and shopping context to deliver relevant offers, airlines are falling behind in this regard with their static content and one-size-fits-all retailing approach. To meet the growing expectations of their customers, the airline industry has expressed a vision for *dynamic offer creation*, which will allow airlines to dynamically bundle and price a set of offers that is customized to the context of the shopping request. Realizing this vision requires significant advancements in both distribution and science. On the distribution side, these advancements will come with the adoption of the New Distribution Capability. On the science side, which is the focus of this paper, little progress has been made despite years of research. In particular, airlines still lack a tractable scientific model to dynamically create and price offer sets at scale. In this paper, we present a novel approach to solve the airline dynamic offer creation problem using a Markov chain choice model. Our model displays attractive qualitative properties—the resulting offers and prices are chosen in such a way as to discourage purchases of unprofitable offers and nudge customers towards more profitable ones. Our model naturally proposes offers that are relevant to the customer, as including irrelevant offers in the offer set leads to a reduction in revenue and ancillary purchase rates. In a simulation study with two customer segments, we find that our model significantly increases ancillary revenue over a naïve, unsegmented pricing model that mimics current state-of-the-art practice. While our studies are conducted under several idealized assumptions, they demonstrate a substantial revenue potential from dynamic offer creation in both unsegmented and segmented applications.

Keywords Dynamic offers · Markov chain choice · Price optimization · Assortment optimization · Airline revenue management · New Distribution Capability

Introduction

The airline retailing experience stands in strong contrast to what customers have come to expect from tech giants such as Apple, Amazon, and Netflix, who have mastered the art

of providing relevant, personalized content to their customers. For example, even though customers often have specific preferences for ancillary services (e.g., seat assignments, checked bags, lounge access, etc.), airlines lack the ability to create and distribute customized offers that are relevant to the traveler making the request.

Consequently, the airline industry has formulated a vision for *dynamic offer creation* (Touraine and Coles 2018) that would enable the industry to move the retailing experience closer to those of sophisticated online retailers. A *dynamic offer* is a combination of products and services shown at a single price provided in real time, on a one-time basis, in response to a specific customer request. In general, a dynamic offer can consist of either a single product (*a la carte offer*) or a bundle of products (*bundled offer*). In this paper, we use the term *offer* to represent both types of dynamic offers.

✉ Kevin K. Wang
kevini@alum.mit.edu

Michael D. Wittman
michael.wittman@amadeus.com

Thomas Fiig
thomas.fiig@amadeus.com

¹ Massachusetts Institute of Technology, International Center for Air Transportation, 77 Massachusetts Avenue, Building 35-217, Cambridge, MA 02139, USA

² Amadeus IT Group, Hedegaardsvej 88, 2300 Copenhagen, Denmark



The airline industry faces two challenges in creating dynamic offers: distribution and scientific modeling. To address the first challenge, the New Distribution Capability (NDC) was developed to remove limitations in the legacy IT systems that prevent airlines from distributing dynamic offers. NDC is a communication standard that allows airlines to generate and distribute offers that are contextual and customized to the traveler making the request. Unlike legacy distribution standards, offers distributed via NDC can also be priced without reference to a predetermined filed fare. We refer readers to IATA (2021) for more information about how NDC will enable distribution of dynamic offers in both direct and indirect distribution channels.

In this paper, we focus our attention on the second challenge: the science of creating and pricing the set of offers from which the customer can choose—the so-called *offer set*. Determining the offer set requires solving several inter-related dynamic pricing and assortment optimization problems: (a) creating relevant offers targeted to the customer’s preferences; (b) pricing the offers while taking the customer choice behavior into account; and (c) creating an assortment of offers that maximizes the airline’s revenue from the shopping session.

Early discussions of offer set creation date back to seminal “bundle pricing” papers by Adams and Yellen (1976) and Schmalensee (1984), who considered a case of two products (goods) that could be either sold separately, offered only as a bundle, or offered both separately and as a bundle. In this simple scenario, the optimal pricing and offer set could easily be determined based on the customer valuations for the two goods. However, in the general case with more than two products, dynamic offer creation becomes a difficult optimization problem.

The complexity of the problem arises from the fact that the purchase probability of an offer depends not only on the price of that offer alone, but also on the other offers in the offer set. This is because customers evaluate each offer relative to the alternatives and make a tradeoff among the perceived utility of the offers and their prices. As a result, optimal pricing requires a joint optimization of the prices of all offers in the offer set.

Many popular discrete choice models used for price optimization, such as the multinomial logit (MNL) and probit, either do not produce logical prices or are mathematically intractable for joint optimization in general. Previous approaches to bundling in the airline industry have been limited to two products (Ødegaard and Wilson 2016; Wang et al. 2021), approximate solutions (Ratliff and Gallego, 2013), or machine-learning approaches (Ettl et al. 2020; Wang 2021). While quantitative evidence of bundle pricing is often seen in other industries, such as for theater tickets (Chu et al. 2011), online music (Elberse 2010), and mobile phone plans (Calzada and Martinez-Santos 2016), these

industries face the same issues and have not found more attractive solutions to the bundle pricing problem in general.

In this paper, we present a novel approach to offer pricing using the Markov chain choice model (MCCM), which we argue is particularly well-suited for dynamic offer creation. Specifically, the structure of the MCCM allows for the decomposition of the multi-dimensional optimization problem into a series of univariate optimization problems that can be solved efficiently, thereby providing the required scalability even for large offer sets.

We will focus specifically on ancillary services in this paper. While dynamic offers in general refer to combinations of both flight itineraries and ancillary services (e.g., a fare family), the ancillary bundling problem is still quite difficult even without adding the additional layers of complexity associated with flight itinerary pricing that would entail seasonality and capacity constraints. However, our framework is entirely general and can also be extended to include flight itinerary pricing if desired (Wang 2022).

We find that the offers produced by the MCCM display intriguing properties. Even though the univariate MCCM price optimizations are not directly aware of customer substitution between offers in the offer set, the resulting offers and prices are chosen in such a way as to discourage purchases of unprofitable offers and nudge customers towards more profitable ones. We also find that the model naturally proposes offers that are relevant to the customer, as including irrelevant offers in the offer set leads to a reduction in revenue and ancillary purchase rates.

Our simulations show that dynamic offer creation using our model can significantly increase ancillary revenue over a naïve, unsegmented pricing model that mimics current state-of-the-art practice. We believe that our model represents a potential breakthrough in offer optimization that would enable dynamic offer creation to be practiced at scale in the airline industry.

Contributions

In this paper, we demonstrate how the Markov chain choice model can be used to solve the airline dynamic offer creation problem at scale. We make several enhancements to the MCCM to adapt it to our setting, including the concept of ancillary relevance, which penalizes the model for showing offers that are not of interest to the customer. We show that our model possesses desirable properties—it is simple, intuitive, and produces logical prices for bundled offers. Moreover, its structure allows for efficient optimization of offer prices, even for large offer sets. Finally, we present a series of simulation results that enable practitioners to understand the qualitative properties and revenue performance of our model.



Organization of the paper

The remainder of this paper is structured as follows. We first describe the business context of airline retailing, followed by a literature review. We then formulate the dynamic offer creation problem and describe how the MCCM represents customer choice among offers in an offer set. Next, we evaluate the performance of our model with a focus on the qualitative aspects. We also assess the relative revenue benefits of segmentation. Finally, we provide conclusions and outline future research directions.

Business context of dynamic offer creation

In this section, we describe the current business context regarding airline retailing and the move towards dynamic offer creation. We first describe how airlines retail today and the limitations of this approach. We then discuss the drivers behind the industry shift towards more modern, customer-centric retailing, and how this shift has inspired new advancements in distribution technology to overcome roadblocks that exist today.

How do airlines retail their offers today?

In the 2000s and early 2010s, retailing strategies differed significantly between airlines with different commercial business models. For example, full-service carriers (FSCs) were traditionally known to offer bundled fares in which ancillary services were included along with the flight. Low-cost carriers (LCCs), on the other hand, were pioneers of “unbundling” retailing strategies, where ancillary services were no longer included with the flight but could instead be purchased optionally. This shift meant that ancillary revenues as a share of total revenue grew significantly. For some low-cost carriers such as Allegiant, Spirit, and Wizz Air, ancillary services now generate more revenue than flights (Sorensen 2021).

Retailing strategies have started to converge between business models. To match the low base fares of LCCs, FSCs introduced new unbundled offers where ancillary services were not included with the flight, whereas some LCCs started to bundle fare products to match the product offerings of full-service carriers (CAPA 2018). The result was a widespread “hybrid” strategy in which customers were first presented with a selection of fare families and branded fares, and then offered additional ancillary services later in the booking flow (Vinod et al. 2018).

Today, airlines rely on a two-step process to retail their products in the indirect distribution channels, such as online travel agencies and metasearch platforms. First, the airline publishes fares with fare distributors such as ATPCO or

SITA; second, the availability computation within the airline’s inventory system determines which filed fares are made available for sale. This allows prices to be pre-computed offline and reduces the real-time process to a lightweight availability computation. This approach is highly scalable and efficient, but also leads to several limitations.

First, because prices for each fare class are computed ahead of time, no contextual or customer-specific information can be used to create customized offers. There is also no ability to dynamically create bundles of ancillary services in the indirect channel—only static, pre-constructed offers can be displayed. Finally, ancillary service purchases are often stored in electronic miscellaneous documents (EMDs) that are attached to the passenger name record (PNR)—this separate storage makes analyzing and tracking ancillary purchases difficult, particularly in real time.

Further, the airlines have limited control over the offer creation in the indirect channel. This is because intermediaries such as the Global Distribution Systems build offers on behalf of the airlines. This can lead to a mismatch of incentives—for example, an airline may wish to display a certain connection or a bundled offer to maximize revenue, whereas a travel agency may bias the display to favor the cheapest, no-frills offer to try to maximize conversion.

It is worth noting that on the direct distribution channels, such as the airline’s own website, airlines have full control over the creation of the offer. All information about the travel context and customer is available to the airline, allowing for the creation of fully customized content. However, the limitations of the indirect channel restrict the airline’s ability to fully utilize these capabilities, due to the desire to maintain price and content consistency between channels. Inconsistency is confusing for customers and creates mistrust, while also causing “channel shift” to the distribution channel that provides the cheapest or most-desired offer.

How do airlines wish to retail their offers in the future?

To accomplish their vision of dynamic offer creation, airlines wish to propose an offer set that is customized based on the customer request and travel context and is consistent across channels. Figure 1 compares airline retailing under the current distribution paradigm (left panel) with dynamic offer creation (right panel). As we focus only on the selection of ancillary services in Fig. 1, this figure can be interpreted as the portion of the shopping flow after the flight itinerary has been selected.

In the traditional shopping flow, customers are presented with shopping pages containing static offers. Each shopping page is displayed sequentially, and the offers (both product and price components) and order of the shopping pages do not change from customer to customer. The prices of the



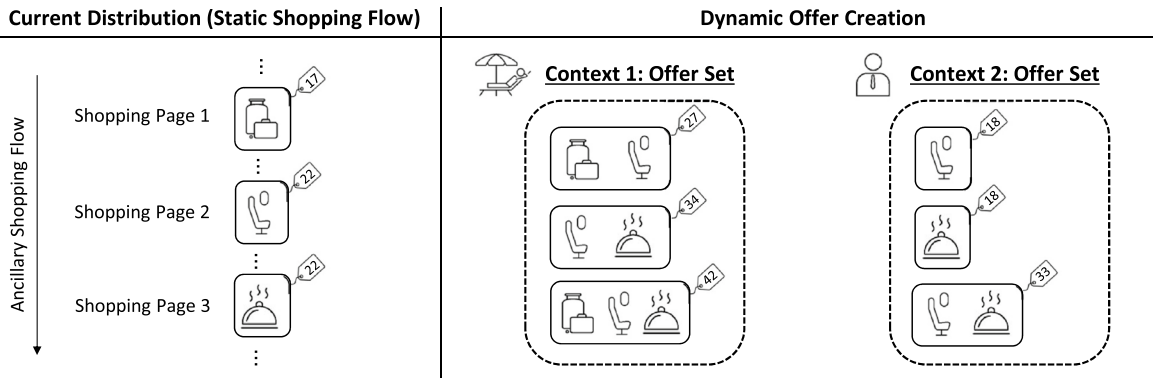


Fig. 1 Ancillary retailing under the current distribution paradigm (left panel) compared to dynamic offer creation (right panel)

products are static and do not vary based on the context of the shopping request nor on the products selected by customers on previous pages. Ideally, the airline would instead want to bundle complementary products together to encourage upsell, and to show only the relevant products to customers on each shopping page.

Under dynamic offer creation, the composition and pricing of offers in the offer set can vary contextually. In this example, the leisure travel context (context 1) is shown a different selection of offers than the business travel context (context 2) for the same itinerary. The price of an offer can also depend on the offer set in which it is contained, as well as the context in which the shopping request is made. In Fig. 1 the price of the bundled offer containing the seat and the meal is \$34 in the offer set returned to context 1 and \$33 in the offer set returned to context 2. These price differences could be larger if there is a significant difference in valuation between the travel contexts.

Dynamic offer creation allows the airline to display the offers that are the most relevant to each context. In Fig. 1, the business travel context is not shown an offer containing the bag, as this ancillary is not as relevant to this context. It also allows the airline to upsell customers to bundled offers by removing some offers from the offer set or pricing them unattractively. For example, a customer in context 1 wishing to purchase the seat must purchase it in a bundle with either the bag or the meal, whereas a customer in context 2 is able to purchase the seat alone.

Industry progress towards the vision of dynamic offer creation

Real-world airline progress towards dynamic offer creation can be mapped through IATA’s Dynamic Offer Capability Model Matrix (Touraine 2021). The matrix is reproduced in Fig. 2, to which we have added annotations.

Each axis of the matrix shows an increasing level of sophistication and dynamicity in either price (horizontal

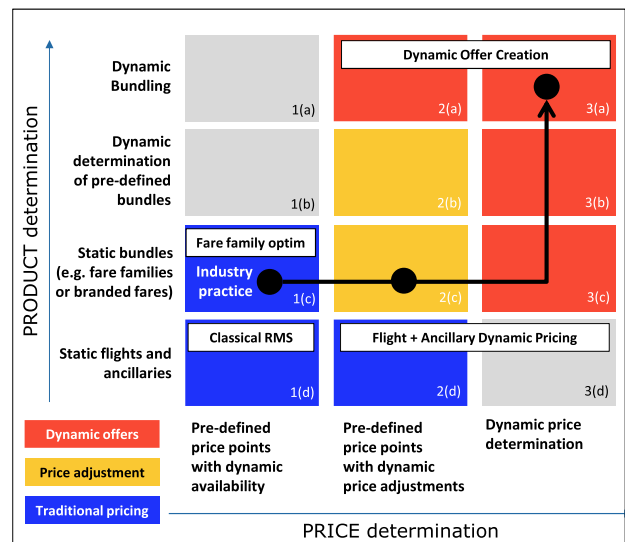


Fig. 2 IATA Dynamic Offer Capability Model Matrix (Touraine 2021, reproduced with permission) with annotations added

axis) or product determination (vertical axis). The lower left corner of the matrix (blue boxes) shows the current distribution framework of statically created offers at pre-defined price points. The top right of the matrix represents dynamic offer creation, which incorporates both dynamic bundling and dynamic pricing of offers without reference to a pre-defined price.

In Fig. 2, the current state of airline offer optimization is represented by the bottom-left corner of the matrix. Classical revenue management systems (RMS) dynamically select which pre-filed price points to make available for static filed products (dynamic availability; box 1(d)).

In the past decades, classical RMS has evolved along the product dimension to optimize static bundles of flights and ancillaries (fare family optimization; box 1(c)) and along the price dimension to dynamically adjust prices from pre-filed price points (dynamic pricing; box 2(d)). The combination



of dynamic bundling and dynamic pricing (dynamic offer creation, box 3(a)) is a significant shift from what is practiced in the industry today.

This gap between current industry practice and the full vision of dynamic offer creation is a result of both inadequacies in legacy distribution and scientific modeling. On the distribution side, the New Distribution Capability (NDC) was introduced to address the technical issues that stem from current distribution standards. NDC is a communication standard that allows airlines to create and distribute dynamic offers in both direct and indirect channels. It allows customer context to be included as part of the offer request, and for offers to be priced without reference to a filed fare. NDC, therefore, enables dynamic offer creation from a technical point of view.

On the science side, dynamic offer creation requires science that enables the airline to create and price a set of offers to optimize one or more airline objectives. As we will describe, customer choice models commonly used in the airline literature do not possess the properties needed to solve the dynamic offer creation problem logically and at scale.

Finding an operational scientific model for dynamic offer creation will be the focus of the remainder of this paper.

Literature review

We will continue to use the matrix from Fig. 2 as a framework to orient our literature review. In Fig. 3, we further annotate the matrix with a list of representative research papers for selected combinations of price and product optimization complexity.

Decades of airline revenue management literature focused on solving the bottom-left box of the capability matrix by optimizing pre-defined products at pre-defined price points. For example, Belobaba (1987) introduced the canonical EMSRb model for airline revenue management assuming independent demand, and Talluri and van Ryzin (2004) showed how the single-leg revenue management problem could be extended to incorporate customer choice in situations with dependent demand. Fiig et al. (2010) showed how the marginal revenue transformation could be used to transform dependent demand to an equivalent independent demand representation, allowing the use of independent demand optimizers in situations with dependent demand.

Moving one box higher in the matrix, the classic RMS problem was expanded to consider choice behavior across a static set of fare families or branded fares. Fiig et al. (2012) and Walczak and Kambour (2014) showed how customer choice across fare families could be incorporated in the marginal revenue transformation, and Hjorth et al. (2018) incorporated overbooking into fare family optimization. In each of these papers, available prices are selected from a set of pre-defined price points for each product.

Continuing to move up in the matrix to the selection of pre-defined, pre-priced bundles, Ratliff and Gallego (2013) described how branded fare bundles can be designed and priced using willingness-to-pay estimates obtained from conjoint analysis on survey data. Their method for constructing the bundle offer involved a multi-dimensional optimization problem that was so complex that an approximate solution had to be found. Their paper focuses on the creation of static branded fare pricing ladders, and not on the dynamic real-time creation and pricing of offers.

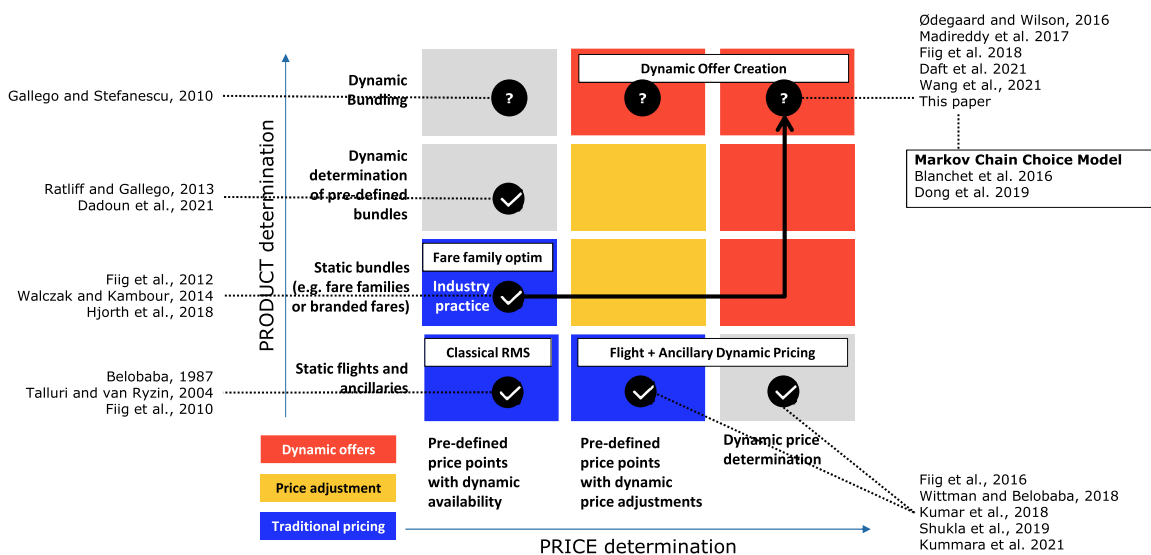


Fig. 3 Annotated IATA Capability Matrix from Dynamic Offer Creation (adapted from Touraine 2021)



Dadoun et al. (2021) described how recommender systems can be used to proactively propose a set of products or services to customers as a function of the items they have previously purchased either in the same shopping session or in past transactions. They describe how many classic algorithms in the recommender system literature can be adapted to the airline revenue management setting. However, they do not consider pricing of the offers.

Finally, there is limited literature focused on dynamic bundling within the airline industry. One exception is Gallego and Stefanescu (2010), who discussed how to construct customized bundles for multiple customer segments based on each segment's valuation for each product. Their model constructs bundle prices as the sum of the prices of each of the bundle's constituents; bundle discounts are not considered.

Significantly more attention has been paid to the dynamic pricing of individual flights or ancillary services, shown in the bottom-right corner of the matrix. Fiig et al. (2016) described a method for performing dynamic price adjustment to flight itinerary prices based on context of the shopping session, and Wittman and Belobaba (2018) investigated the revenue benefits of incremented and discounted dynamic price adjustments. Kumar et al. (2018) showed how an active learning dynamic pricing method could increase revenue while estimating the demand behavior of customers. Shukla et al. (2019) and Kummara et al. (2021) proposed machine-learning approaches for dynamic pricing of ancillary services based on observable customer characteristics. These papers focus exclusively on the pricing portion of the dynamic offer creation problem, and do not consider the selection of products to display to customers.

The top-right region of the matrix associated with dynamic offer creation has been much discussed, but little scientific progress has been made. Several papers have described the technical frameworks that would be needed to enable dynamic offer set selection and pricing. These papers are often quite conceptual in nature with limited focus on the specific models used in dynamic offer creation. For example, Madireddy et al. (2017) described a piecewise approach for offer optimization consisting of separate modules that estimate customer willingness-to-pay, perform segmentation, and generate sets of possible offers, based on the information available from RMS and e-commerce platforms. Fiig et al. (2018) described an end-to-end flow to dynamically price flights and ancillary services and discuss how this system would interact with current distribution practices. Daft et al. (2021) and Schubert et al. (2021) proposed alternative flows and discussed frameworks for airlines to move towards dynamic offer creation.

The dynamic offer creation problem is difficult to solve in particular for larger bundle sizes and in offer sets with multiple bundles. One of the challenges is the incompatibility

of popular discrete choice models, such as the probit and the multinomial logit (MNL), with the bundle pricing problem. Aydin and Ryan (2000) and Gallego and Wang (2014), among others, have proven that under an MNL, the optimal markup over the product cost is identical for all jointly optimized products. This poses a particular problem in the dynamic offer context. For example, if all products had zero marginal cost of provision, all offers would have the same optimal price under the MNL, whether or not they are bundled. This leads to illogical pricing behavior, since an offer containing only one ancillary would be priced the same as a bundled offer containing many ancillaries.

Papers that propose specific models for dynamic offer creation typically focus on a minimal bundle of a flight itinerary with a single ancillary service. Ødegaard and Wilson (2016) considered a flight-plus-ancillary dynamic pricing problem, where the seller aims to maximize total expected revenue across the two products. Bundle discounts and mixed bundling are not explicitly considered. Wang et al. (2021) solved a flight-plus-ancillary bundle pricing problem with a focus on optimal bundle pricing and composition. While their probit-based customer choice model can in principle be extended to multiple ancillary services, the complexity increases significantly.

Ettl et al. (2020) considered a case with 11 ancillary services where a customer is shown an offer set of individual services and a single, customized bundle, which can be offered at a discount from the myopic prices. They find that personalized bundling and pricing can increase revenue by 2%–7% compared to a fixed-price base case without bundling. They do not consider the case of multiple bundles in an offer set.

In this paper, we will focus our attention on solving the dynamic offer creation problem by applying the Markov Chain Choice Model (MCCM), a relatively new discrete choice model.

Blanchet et al. (2016) introduced the MCCM and described how it can exactly replicate several popular random utility models (RUMs) such as the MNL and the generalized attraction model (GAM). It has been speculated that the MCCM could in fact replicate *any* RUM (Berbeglia 2016). However, this was recently proven wrong by Ma (2021) who demonstrated that an MCCM cannot in general replicate a nested logit model.

Although MCCM cannot replicate every RUM exactly, Blanchet et al. (2016) demonstrated that under mild assumptions, the Markov chain model is a good approximation for any RUM. Specifically, they determine theoretical bounds and demonstrate numerically for selected scenarios that the average maximum relative error of the choice probabilities from the approximation is less than 3% across the offer set.

Dong et al. (2019) extended the MCCM framework to explicitly incorporate the prices of products in the choice



set. They point out that the optimal prices for an assortment under the MCCM can be solved using a series of univariate price optimizations. This approach provides significant benefits compared to other discrete choice models such as the probit model, which can be difficult to scale effectively to many products. Despite these advantages, no paper to date has investigated the use of MCCM in a dynamic bundling and dynamic pricing problem such as dynamic offer creation. We aim to fill this gap in the literature.

Dynamic offer creation problem formulation

Definition of offer

An *offer* consists of two parts: a *set of products* and a *price*—as an unbreakable entity. However, industry convention uses the term *offer* to mean either the set of products only (which we call an *unpriced offer*) or the whole entity (which we call a *priced offer*).

Suppose an airline has a catalog of K atomic ancillary services $\mathcal{A} = \{1, 2, \dots, K\}$, where we let $a = 1, \dots, K$ denote the individual ancillaries. We define an *unpriced offer* $O \subseteq \mathcal{A}$ as a subset of one or more ancillaries. To create the universe of all possible offers, we consider the powerset $\Omega = \mathcal{P}(\mathcal{A}) \setminus \{\emptyset\}$ excluding the empty set (no ancillaries purchased). The cardinality is $|\Omega| = 2^K - 1$.

Definition of offer set

We define an *unpriced offer set* $S = \{O_1, \dots, O_N\}$, $S \subseteq \Omega$ as a collection of one or more offers that are displayed to the customer. The cardinality is denoted by $N = |S|$ and satisfies $1 \leq N \leq |\Omega|$. We can generate all non-empty offer sets $S \subseteq \Omega$ by considering the powerset $\mathcal{P}(\Omega) \setminus \{\emptyset\}$, which has a cardinality of $2^{2^K - 1} - 1$. Note that the customer always has the choice of not purchasing any offer, so every offer set always includes a no-purchase option $O_0 = \emptyset$. Hence, when we use the notation S it should tacitly be understood as the union $S \cup \{O_0\} = \{O_0, O_1, \dots, O_N\}$.

The *priced offer set* is constructed by associating an unpriced offer set S with a vector of prices of each offer in the offer set: $\mathbf{p}(S) = (p_1(S), \dots, p_N(S))$. We will also allow the notation $p_i(S) = p_i = p_{O_i}$ that omits the offer set S when it is clear from the context, as well as explicitly stating the offer O_i as an index. The price of the no-purchase option is always $p_0(S) = 0$.

Example

Consider a case with $K = 3$ ancillary services $\mathcal{A} = \{1, 2, 3\}$, for example representing a *bag*, *seat*, and a *meal*. Then $\Omega = \mathcal{P}(\mathcal{A}) \setminus \{\emptyset\}$, is constructed by the $|\Omega| = 7$ offers:

$$\begin{aligned} \Omega &= \{O_1, O_2, O_3, O_4, O_5, O_6, O_7\} \\ &= \{\{1\}, \{2\}, \{3\}, \{1, 2\}, \{1, 3\}, \{2, 3\}, \{1, 2, 3\}\}. \end{aligned} \quad (1)$$

The ordering above is constructed by interpreting the ancillary services index as digits, e.g., $\{2, 3\} = 23$, that are then ordered in ascending order—so-called “Banker’s ordering” (Loughry et al. 2000).

An offer set S is a subset of Ω , $S \subseteq \Omega$. Examples include $S_1 = \{\{1\}, \{2\}, \{1, 2, 3\}\}$ and $S_2 = \{\{1\}, \{3\}, \{1, 2, 3\}\}$. If $S = \Omega$, then we denote S the *full offer set*.

Definition of optimization objective

Given a priced offer set, we use a discrete choice model $\mathcal{P}(O_i|S, \mathbf{p}(S))$ to specify the probability that a customer chooses offer O_i among the alternatives in the offer set $S = \{O_1, \dots, O_N\}$.

From the framework of discrete customer choice models, offers within an choice set must be *finite*, *collectively exhaustive* and *mutually exclusive*, meaning that a customer makes *one and one choice only* among a finite set of alternatives. In practice, an airline could display an offer set from which the customer could make *multiple* selections. However, we can always map the case of a multiple selection offer set into an equivalent single selection choice set. For example, suppose an airline displays the offer set $\{\{1\}, \{2\}\}$, allowing the customer to make multiple selections. This can be mapped to an equivalent single selection choice set $\{\{1\}, \{2\}, \{1, 2\}\}$. Thus, without loss of generality we will assume that the offer set is a “proper” single selection choice set.

The *expected net revenue* of an offer O_i in offer set S is computed by multiplying the offer’s *contribution*—that is, its price p_i minus the unit cost of providing the offer c_i —by the purchase probability $\mathcal{P}(O_i|S, \mathbf{p}(S))$ that a customer will choose that offer from the offer set. We follow revenue management industry convention and simply use the term *revenue* to tacitly mean net revenue.

We define the expected revenue for the entire offer set, $\pi(S, \mathbf{p}(S))$, as the sum of the expected revenue from all offers in the offer set:

$$\pi(S, \mathbf{p}(S)) = \sum_{i=1}^N (p_i - c_i) \mathcal{P}(O_i|S, \mathbf{p}(S)). \quad (2)$$

The *dynamic offer creation problem* can then be separated into the following two optimization subproblems:

- *Offer Set Pricing*: For a given unpriced offer set S , determine the optimal prices: $\mathbf{p}^*(S) = \operatorname{argmax}_{\mathbf{p}} \pi(S, \mathbf{p}(S))$
- *Offer Set Selection*: For given priced offer sets $\mathbf{p}(S)$, $S \subseteq \omega$, $\omega \subseteq \Omega$, determine the revenue maximizing offer set $S^* = \operatorname{argmax}_{S \subseteq \omega} \pi(S, \mathbf{p}(S))$



The *joint offer set pricing and offer set selection* problem does not add new complexity, given we can perform an exhaustive search for the optimal offer set $S^* \subseteq \omega$ among the priced offer sets.

The Markov chain choice model for dynamic offer creation

Alptekinoglu and Semple (2016) suggest several properties that are considered *desirable* in a discrete choice model for pricing and assortment optimization problems. They state that a discrete choice model $\mathcal{P}(O_i|S, \mathbf{p}(S))$ should be:

- *Scalable*: purchase probabilities can be computed even for large offer sets.
- *Customizable*: the offer set can be tailored to the individual customer.
- *Interpretable*: a parametric discrete choice model—allowing interpretation.
- *Estimable*: the model parameters can be estimated from training data.
- *Tractable*: the revenue maximizing prices $\mathbf{p}^*(S)$ or offer set S^* can easily be determined, even for large offer sets.

In our setting, we require additionally that the revenue maximizing prices $\mathbf{p}^*(S)$ satisfies:

- *Logical prices*: the prices $\mathbf{p}^*(S)$ must satisfy logical restrictions with respect to the prices of bundles and

individual components. For offer set $\{\{1\}, \{2\}, \{1, 2\}\}$, $\max(p_{\{1\}}, p_{\{2\}}) \leq p_{\{1,2\}}$ and $\max(p_{\{1\}}, p_{\{2\}}) \leq p_{\{1\}} + p_{\{2\}}$.

As we have discussed in the literature review, prominent discrete choice models, such as the multinomial logit, nested logit, mixed logit, or probit, do not satisfy these desirable properties. Even for the analytically simple multinomial logit model (MNL), Hanson and Martin (1990) showed that the *offer set pricing* problem is generally non-concave or quasi-concave in prices, resulting in a model that is not tractable for large offer sets. The probit model is neither scalable nor tractable, since computing optimal prices requires evaluating convolution integrals that become more complex as the choice set grows.

Instead, we consider the Markov chain choice model (MCCM) for the dynamic offer creation problem. As we will demonstrate in the simulation section, the MCCM satisfies all the desirable properties above.

The Markov chain choice model

Consider a priced offer set $S = \{O_1, \dots, O_N\}$ with associated prices $\mathbf{p}(S)$. The Markov chain choice model represents a customer’s random walk between individual offers in the offer set. The customer exits the Markov chain by either purchasing one of the offers or electing to leave the shopping flow without purchasing any offer. The MCCM as seen from the customer perspective is shown in the left panel of Fig. 4.

The Markov chain is constructed in the following way: first, a customer arrives at a random offer $O_i \in S$ with

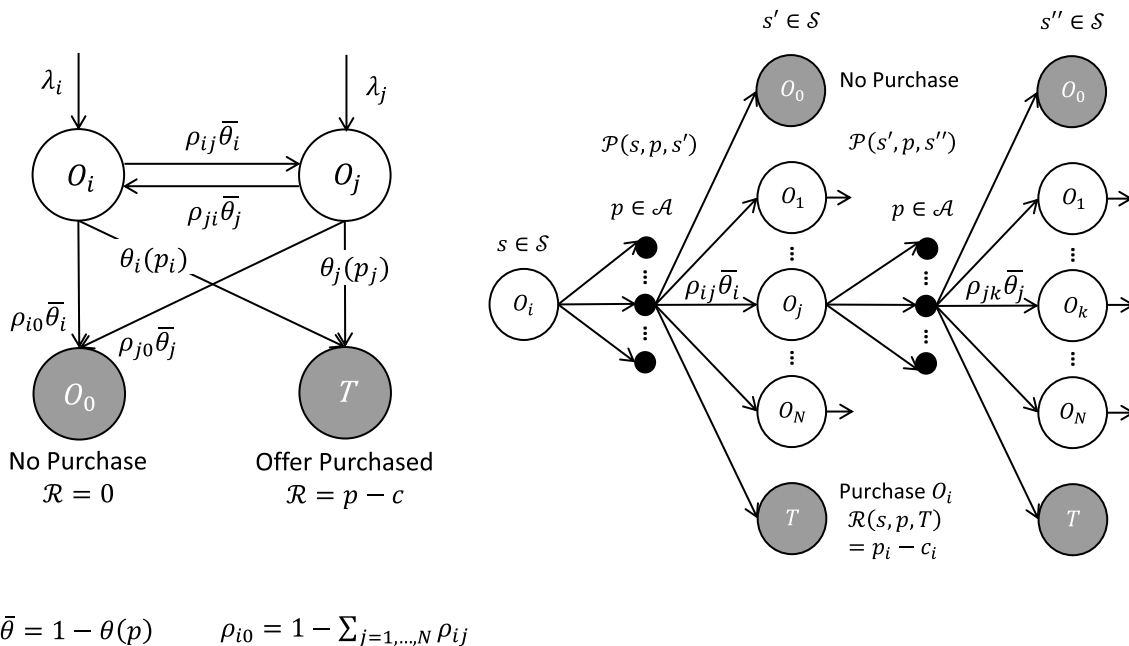


Fig. 4 The Markov chain choice model seen from the customer perspective (left panel) and the price optimization perspective (right panel)



probability λ_i , $i = 1, \dots, N$, which denotes the first-look probability of a customer entering the shopping session. Upon visiting an offer $O_i \in S$, the customer makes one of three possible transitions:

- Purchase offer O_i with the so-called *myopic* purchase probability $\theta_i(p_i)$, since the probability only depends on p_i but not on the other prices in the offer set. The customer then exits the chain (transitions to the terminal state T).
- Transition to look at another offer $O_j \in S$, $j = 1, \dots, N$ with probability $\rho_{ij}(1 - \theta_i(p_i))$, where ρ_{ij} is a memory-less transition rate. Self-transition is not allowed; thus, $\rho_{ii} = 0$, $\forall i = 1, \dots, N$.
- Transition to the terminal no-purchase state $O_0 = \emptyset$. This occurs with a probability $(1 - \theta_i(p_i))\rho_{i0} > 0$. It represents a customer who leaves the shopping flow without purchasing anything.

Note that the two exit options—no-purchase $O_0 = \emptyset$, and purchase T —are absorbing states. Hence $\rho_{0T} = \rho_{T0} = 0$ and $\rho_{0j} = \rho_{Tj} = 0$ for $j = 1, \dots, N$.

It may seem artificial to imagine customers randomly transitioning between offers. However, it is important to highlight that MCCM does not have to directly replicate how individual customers evaluate and purchase offers. Rather, through appropriate choices of the parameters λ_i , $\theta_i(p_i)$, and ρ_{ij} , the MCCM can approximate collective behavior from an ensemble of customers for many types of random utility models, as shown by Blanchet et al. (2016) and others.

One way of thinking about the MCCM is as a mathematical “trick” that allows us to transform a RUM to an approximate MCCM. As we will show below and in our simulation section, a formulation based on MCCM allows an efficient solution of the dynamic offer creation problem, even though the original problem is difficult to solve.

In this paper, we will assume that the MCCM accurately reflects the customer choice behavior, whereas in practice, the model parameters need to be estimated and fit to the observed customer choice behavior. Simsek and Topaloglu (2018) demonstrated that maximum likelihood estimation and expectation–maximization can be used to estimate the MCCM parameters in principle. Berbeglia et al. (2021) showed that the MCCM requires more data and computational time to estimate than other discrete choice models due to its large number of parameters, but that once estimated, the MCCM fits the underlying data better. They showed that the estimation can be sped up by constraining the model and reducing the number of parameters to be estimated, but that the constrained MCCM no longer consistently fits the customer choice data better than other discrete choice models. However, the extent to which this degradation in model fit

impacts the performance of the MCCM price optimization was not investigated.

With offer set selection, the complexity increases even further, as each offer set potentially has its own set of MCCM parameters to be estimated. In practice, airlines would not have sufficient historical data with different offer sets at different price points to estimate multiple MCCMs. Gupta and Hsu (2020) demonstrated that it is possible to recover MCCM parameters for a large offer set from known parameters for smaller subsets, but further research is needed into the practicality of converting parameters between offer sets and on the estimation of MCCM parameters in general. Wang (2021) found that price experimentation conducted with Thompson sampling could be a promising means to reduce the cost of collecting the historical data, and thus, the cost of estimating the MCCM.

Price optimization for the Markov chain choice model

The structure of the MCCM allows the price optimization to be represented as a Markov Decision Process (MDP). Formally, an MDP is defined as a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{P})$ where \mathcal{S} is the state space, \mathcal{A} is the action space, $\mathcal{R} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ is the reward function which takes a finite number of values corresponding to the price points of the offers, and $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is the state transition probability function. For the MCCM setting, note that $\mathcal{A}, \mathcal{R}, \mathcal{P}$ are all time independent which simplifies the formulation.

We can define the state space $\mathcal{S} = S \cup \{O_0, T\}$ as the union of the offer set $S = \{O_1, \dots, O_N\} \subseteq \Omega$, and the two absorbing states $\{O_0, T\}$. The action space $\mathcal{A} = \mathbb{R}_+$ is defined as the set of possible price points. The state transition process of the MCCM is represented in the right panel of Fig. 4. The system starts in state $s = O_i$, $i = 1, \dots, N$; the agent (RMS) selects the action (price) $p \in \mathcal{A}$ according to a deterministic policy $\mu : \mathcal{S} \rightarrow \mathcal{A}$. The system then enters its successor state s' with probability $\mathcal{P}(s, p, s')$ while the agent receives a reward $\mathcal{R}(s, p, s')$. The process is repeated until termination—either a customer purchasing an offer or leaving the chain.

To formulate the MDP we need to specify the transition probabilities $\mathcal{P}(s, p, s')$ and rewards $\mathcal{R}(s, p, s')$ based on the MCCM. We obtain for $i = 1, \dots, N$; $j = 1, \dots, N$:

$$\mathcal{P}(O_i, p_i, s') = \begin{cases} \theta_i(p_i) & s' = T \\ (1 - \theta_i(p_i))\rho_{ij} & s' = O_j \\ (1 - \theta_i(p_i))(1 - \sum_{j>0}\rho_{ij}) & s' = O_0 \end{cases} \quad (3)$$

Similarly, the rewards become



$$\mathcal{R}(O_i, p_i, s') = \begin{cases} p_i - c_i & s' = T \\ 0 & s' = O_j \\ 0 & s' = O_0 \end{cases} \quad (4)$$

For the absorbing states we have trivial conditions $\mathcal{P}(O_0, p, s') = \mathcal{P}(T, p, s') = 0$, and similarly for the rewards $\mathcal{R}(O_0, p, s') = \mathcal{R}(T, p, s') = 0$, for all $s' \in \mathcal{S}$.

We introduce the *value function* $V^*(s)$ as the expected revenue from the revenue maximizing policy μ^* : $\mathcal{S} \rightarrow \mathcal{A}$ in a state (offer) $s = O_i \in \mathcal{S}$. The value function satisfies the Bellman equation:

$$V^*(s) = \max_p \sum_{s' \in \mathcal{S}} \mathcal{P}(s, p, s') [\mathcal{R}(s, p, s') + V^*(s')]. \quad (5)$$

The Bellman equation can be solved iteratively,

$$V_{k+1}(s) = \max_p \sum_{s' \in \mathcal{S}} \mathcal{P}(s, p, s') [\mathcal{R}(s, p, s') + V_k(s')], \quad (6)$$

where we have introduced an iteration index $k = 1, 2, 3, \dots$. This is guaranteed to converge to the optimal solution from any starting point $V_k(s) \rightarrow V^*(s)$, for $k \rightarrow \infty$, because the Bellman operator is a contraction mapping (Blackwell 1965). Thus:

$$p^*(S) = \operatorname{argmax}_p V^*(s).$$

It is important to state that the convergence is quite general and the application to the Markov chain choice model cf. Theorem 3 in Dong et al. (2019) is just a special case. The iterative improvements can be performed along one price dimension at a time, greatly reducing the complexity of the optimization problem. The Markov chain choice model is particularly attractive since the price dimension is separated into N one-dimensional price dependencies $\theta_i(p_i)$, $i = 1, \dots, N$. Hence, we naturally construct an iteration procedure: for each $i = 1, \dots, N$ optimize p_i in Eq. (6), while keeping the remaining $p_j, j \neq i$ constant. Repeat until convergence.

Determining the purchase probability for the Markov chain choice model

In the previous section we determined the optimal prices that maximize $\pi(S, \mathbf{p}(S))$. This was done without explicitly determining the purchase probabilities $\mathcal{P}(O_i | S, \mathbf{p}(S))$.

However, the purchase probabilities can be computed by observing that the purchase for an offer is a result of counting the number of visits $v_i(\mathbf{p}(S))$ at offer O_i and multiplying by the myopic purchase probability,

$$\mathcal{P}(O_i | S, \mathbf{p}(S)) = v_i(\mathbf{p}(S)) \theta_i(p_i). \quad (7)$$

Under the Markov chain choice model, the expected number of visits $v_i(\mathbf{p}(S))$ at offer O_i may be written as a balance equation at the vertex of offer O_i (Dong et al. 2019):

$$v_i(\mathbf{p}(S)) = \lambda_i + \sum_{j=1, \dots, N} \rho_{ji} (1 - \theta_j(p_j)) v_j(\mathbf{p}(S)). \quad (8)$$

This equation states that the expected number of times that a customer visits offer O_i (LHS) consists of the direct arrival probability of a customer λ_i at offer O_i summed with the probability of customers transitioning from another unpurchased offer O_j to O_i (RHS). Inserting the optimal prices in (8) we obtain an $N \times N$ linear system of equations, which is easily solved for the visits, and finally the visits are used in expression (7) for the purchase probability. This concludes the model derivation.

Specifying the myopic purchase probability for dynamic offer creation

In this section we formulate a random utility model for the myopic purchase probability that is specific to the dynamic offer creation problem. We assume the following:

1. *Customer segments*: We have $l = 1, \dots, L$ customer segments with population weights w_l ($\sum_{l=1}^L w_l = 1$).
2. *Gaussian valuation*: We model the valuation (or willingness-to-pay) for an atomic ancillary using a Normal distribution for each customer segment l : $\varrho_{a,l} \sim N(\mu_{a,l}, \sigma_{a,l}^2)$, similar to Schmalensee (1984) and Fuerderer et al. (1999). Other probability distributions can also be used within the MCCM. For example, using Gamma or log-normal distributions can ensure strictly positive valuations, yet they are cumbersome to work with, because the valuation of the offer computed by summing the constituting ancillaries' standalone valuations (see Additivity below), no longer can be expressed in a simple closed form expression. Some researchers (Chu et al. 2008) assume the property of *free disposal*, where any negative valuation generated for a customer is set equal to zero.
3. *Additivity*: The valuation for a bundled offer O is equal to the sum of the constituting ancillaries' standalone valuations for each segment l . When each ancillary's valuation is independent and Normal distributed, the bundle valuation is also Normal distributed with $\varrho_{O,l} \sim N(\mu_{O,l}, \sigma_{O,l}^2)$, where $\mu_{O,l} = \sum_{a \in O} \mu_{a,l}$ and $\sigma_O^2 = \sum_{a \in O} \sigma_{a,l}^2$. This would not hold in the case of complementary or substitute ancillaries with positive or negatively correlated valuations.
4. *Relevance*: Customers evaluate the relevance of each atomic ancillary as a binary decision, which is distributed according to a Bernoulli distribution, $B(\phi_{a,l})$,



where we allow for segment-specific parameters. Thus, the atomic ancillary a is relevant for customer segment l with probability $\phi_{a,l}$ and irrelevant with probability $1 - \phi_{a,l}$. The relevance of a bundled offer O is assumed to be the product of its ancillary relevancies which then becomes Bernoulli distributed $B(\prod_{a \in O} \phi_{a,l})$.

With the assumptions above we can write the myopic purchase probability for an offer O_i for customer segment l , priced at p_i as:

$$\theta_{i,l}(p) = \left(\prod_{a \in O_i} \phi_{a,l} \right) \left(1 - \Phi \left(\frac{p_i - \mu_{O_i,l}}{\sigma_{O_i,l}} \right) \right), \quad (9)$$

where $\Phi(x)$ is the cumulative distribution function (cdf) of $N(0, 1)$.

When offer prices are optimized separately for each customer segment, independent MCCMs can be formulated for each segment using segment-specific parameters. When price segmentation is not desired, we can construct an unsegmented MCCM from a mixture of multiple customer segments $l = 1, \dots, L$ by computing a weighted average of the segment-specific myopic purchase probability with population weights w_l , ($\sum_{l=1}^L w_l = 1$):

$$\theta_i(p_i) = \sum_{l=1}^L w_l \theta_{i,l}(p_i) \quad (10)$$

Results and performance

In this section, we report on the performance of dynamic offer creation in a hypothetical market with two customer segments. We investigate the qualitative properties of the MCCM to assess if the model exhibits the desirable properties as discussed in the previous section as well as the revenue benefits of MCCM with varying degrees of sophistication.

Experimental setup

For our experiments, we assume that an airline sells three ancillary services $\mathcal{A} = \{1, 2, 3\}$ and wishes to solve the *joint offer set selection and pricing problem* for the full offer set $S = \Omega$ (Eq.1). To have a concrete scenario in mind consider the ancillaries to be: $a_1 =$ “bag,” $a_2 =$ “seat,” and $a_3 =$ “meal.”

The simulation is setup with two customer segments (named segment 1 and segment 2). We assume an equal mix of customers from both segments, hence $w_1 = w_2 = 1/2$. The ancillary valuation μ , σ and relevance ϕ are specified by the segment-specific parameters in Table 1.

Table 1 Default parameters for the Markov chain choice model for the two customer segments

	O_i (Offer)						
	O_1	O_2	O_3	O_4	O_5	O_6	O_7
	{1}	{2}	{3}	{1,2}	{1,3}	{2,3}	{1,2,3}
$\mu_{i,1} = \mu_{i,2}$	\$10	\$20	\$20	\$30	\$30	\$40	\$50
$\sigma_{i,1} = \sigma_{i,2}$	\$3	\$6	\$6	\$6.7	\$6.7	\$8.5	\$9
$\phi_{i,1}$	75%	75%	75%	56%	56%	56%	42%
$\phi_{i,2}$	5%	95%	25%	4.8%	1.3%	23.8%	1.2%
$\lambda_{i,1} = \lambda_{i,2}$				1/ S			
ρ_{ij}				1/ S if $i \neq j$, 0 if $i = j$			
c_i	\$0						
$w_{i,1} = w_{i,2}$	0.5						

- *Segment 1*: Leisure style customer, who finds the three ancillaries equally and highly relevant ($\phi_{i,1} = 75\%$ for $i = 1, 2, 3$). In terms of valuation, ancillaries 2 and 3 ($\mu_2 = \mu_3 = \$20$ and $\sigma_2 = \sigma_3 = \$6$) are valued twice as high as ancillary 1 on average ($\mu_1 = \$10$ and $\sigma_1 = \$3$).
- *Segment 2*: Business style customer, who primarily is interested in ancillary 2, $\phi_{2,2} = 95\%$ while having low relevance for the other two ancillaries ($\phi_{1,2} = 5\%$, $\phi_{3,2} = 25\%$). Those customers that do find an ancillary relevant have the same average valuation as segment 1 for all three ancillaries.

To evaluate the qualitative properties of MCCM, we make simplifying assumptions about its parameters. We assume that both customer segments can be modeled “perfectly” using an MCCM with the parameters shown in Table 1. We assume that the arrival probability λ is identical for all offers in the offer set, hence $\lambda_i = 1/N$ for $i = 1, \dots, N$, with $|S| = N$ offers. Further, the transition probability between offers is assumed to be $\rho_{ij} = 1/N$, for $i = 1, \dots, N$, $j = 1, \dots, N$ and $j \neq i$ ($\rho_{ii} = 0$ otherwise). Note that as we change the offer set, e.g., by removing an offer O , the cardinality of the reduced offer set becomes $|S \setminus \{O\}| = N - 1$. Thus, the arrival and transition probabilities are redistributed among the remaining offers. We have also verified that the MCCM can be solved with non-uniform transition and arrival probabilities and that the optimized prices are reasonable.

Experimental scenarios

We organize our investigations along two axes of interest: Offer Set Selection and Offer Set Pricing:

- *Offer Set Selection*—indicating whether MCCM is used to customize the selection of offers displayed to the customer.



- With *no Offer Set Selection*, the full offer set Ω (consisting of all possible offers) is displayed to all customers.
 - With *Unsegmented Offer Set Selection*, the offer set S^* with the highest expected revenue is displayed ($S^* \subseteq \Omega$), but the offer set does not vary by customer segment.
 - With *Segmented Offer Set Selection*, the displayed offer set $S^* \subseteq \Omega$ may also vary by customer segment.
- *Offer Set Pricing*—indicating whether MCCM is used to optimize prices for each offer in the offer set.
 - *Myopic a la carte pricing*: Each offer is priced independently and without bundle discounts, ignoring substitution among offers.
 - *Unsegmented pricing*: MCCM computes prices for each offer in the offer set. Prices of an offer can vary by offer set, but not by customer segment.
 - *Segmented pricing*: MCCM computes prices for each offer in the offer set. Prices of an offer can vary both by offer set and by customer segment.

Table 2 summarizes the scenarios that are evaluated. Note that in the case of segmentation it is assumed in the simulations that the airline can perfectly identify the customer segments and distribute segmented offer sets.

Baseline pricing—Myopic a la carte pricing

We construct a simple baseline pricing strategy that we call *myopic a la carte pricing* (ALC), which rests on three key assumptions: (a) the airline knows the true customer valuation distribution for each ancillary service; (b) the price for each ancillary is optimized myopically without considering the other offers in the offer set; and (c) no bundle discounts are offered.

The myopic optimal price p_a^* for ancillary $a = 1, \dots, K$, can be determined by maximizing the revenue from the myopic purchase probability.

$$p_a^* = \operatorname{argmax}_{p_a} \left(1 - \Phi \left(\frac{p_a - \mu_a}{\sigma_a} \right) \right) (p_a - c_a). \quad (11)$$

The myopic price of a bundled offer $p(O)$ is simply the sum of the individual myopic prices $p(O) = \sum_{a \in O} p_a^*$.

In Table 3, we show the unsegmented myopic a la carte prices $p^*(\Omega)$, the corresponding purchase rates $\mathcal{P}(\Omega)$ as computed using the MCCM and expected contribution $\pi(\Omega)$ for the two segments.

While the prices are identical for the two customer segments, we observe substantial differences in the purchase behavior. For segment 1, for which all ancillaries are equally

Table 2 List of scenarios with their two-letter codes

Scenario	Offer set selection	Offer set pricing	Location
<i>FM</i>	None (Full Offer Set)	Myopic	Table 3
<i>FU</i>	None (Full Offer Set)	Unseg. MCCM	Table 4, Fig. 5
<i>UU</i>	Unsegmented	Unseg. MCCM	Table 5
<i>FS</i>	None (Full Offer Set)	Seg. MCCM	Table 6
<i>SS</i>	Segmented	Seg. MCCM	Table 6

First letter represents Offer Set Selection (Full, Unsegmented, Segmented). Second letter represents Offer Set Pricing (Myopic, Unsegmented, Segmented)

relevant, we observe a 14.5% purchase rate for offers consisting of a single ancillary, and a 9.2% purchase rate for the bundle of all three ancillaries. For segment 2, we see a higher purchase rate of 37.6% for offer {2}, since this segment sees ancillary 2 as highly relevant compared to the other ancillaries.

We can compute the expected revenue across both customer segments by averaging the expected revenue from each segment $\pi(\Omega) = (\$18.25 + \$11.13)/2 = \$14.69$.

The myopic a la carte pricing represents an optimistic baseline given that airlines in general have limited knowledge of the customers' actual valuation, and that airlines manually file prices to fare aggregators, which would typically be set using business rules at an aggregated level. Further, because we selected a scenario where the customer segments have identical valuations for each ancillary, the optimized prices are myopically optimal for both segments. In reality, different segments will have different ancillary valuations, which would reduce the performance of this unsegmented myopic pricing baseline.

Full offer set, Unsegmented pricing—(FU)

Next, we turn to price optimization based on the MCCM. In this scenario, we restrict the prices to be unsegmented using the myopic purchase probability for two segments (Eq. 10). Here the segments are differentiated only in their relevance parameters, hence $\phi_i = w_1 \phi_{i,1} + w_2 \phi_{i,2}, i = 1, \dots, N$. The results are shown in Table 4.

First, we note that the Unsegmented MCCM prices are higher than the myopic ALC baseline for all offers. This price increase is especially strong for the single ancillary offers. For example, the price of offer {2} increases from \$15.48 under myopic ALC pricing to \$21.71 under MCCM, while the price of the full bundle increases from \$38.69 to \$42.11 under MCCM.

We observe that MCCM exhibits a mixed bundling strategy, where the atomic offers are priced in a way that incentivizes customers to purchase bundles at a discount (Kobayashi 2005). For example, the price of bundle {2, 3} is \$34.54 under MCCM pricing, which represents a 20.6%



Table 3 Unsegmented myopic ALC prices and their purchase rates $\mathcal{P}(\Omega)$ and expected contribution $\pi(\Omega)$ from the two customer segments

Ω	{1}	{2}	{3}	{1,2}	{1,3}	{2,3}	{1,2,3}	$\mathcal{P}(\Omega)$	$\pi(\Omega)$
Price	\$7.74	\$15.48	\$15.48	\$23.22	\$23.22	\$30.95	\$38.69		\$14.69
Seg 1	14.5%	14.5%	14.5%	11.7%	11.7%	11.9%	9.2%	88.0%	\$18.25
Seg 2	1.8%	37.6%	9.2%	1.9%	0.5%	9.7%	0.5%	61.2%	\$11.13

Table 4 Unsegmented MCCM prices and their expected purchase rates and revenues from two segments of customers choosing from the full offer set according to MCCMs with different offer relevance

Ω	{1}	{2}	{3}	{1,2}	{1,3}	{2,3}	{1,2,3}	$\mathcal{P}(\Omega)$	$\pi(\Omega)$
Price	\$17.00	\$21.71	\$21.77	\$27.39	\$27.40	\$34.54	\$42.11		\$16.71
Seg 1	0.3%	10.4%	10.3%	13.2%	13.2%	15.1%	12.3%	74.9%	\$22.21
Seg 2	0.0%	23.8%	6.0%	1.9%	0.5%	11.1%	0.6%	43.9%	\$11.21

savings over purchasing the two ancillaries separately at total price of \$43.48. Such a pricing strategy can be observed in practice at low-cost carriers, which often sell discounted bundles combining services such as seat assignments and checked luggage alongside the atomic ancillaries.

The increase in the prices of the atomic offers and the discount of the bundled offers leads to an increase in customers buying bundles. This can be seen from customer segment 1, where we observe an increased purchase rate for the full bundle {1, 2, 3} from 9.2% with myopic ALC to 12.3% under Unsegmented MCCM. Similarly, the share of customers purchasing a single ancillary decreases from 43.5% to 21.0%. As a result, the total ancillary purchase rate decreases for both segments under Unsegmented MCCM pricing compared to myopic ALC pricing, but the expected revenue per customer increases by 13.8% from \$14.69 to \$16.71 with Unsegmented MCCM.

Sensitivity analysis of Full offer set, Unsegmented pricing—(FU)

In this section, we perform sensitivity tests to see how MCCM prices depend on the input parameters. We deviate from the default parameters in Table 2 by varying one set of parameters at a time, while holding the remaining parameters constant at their default level. To simplify interpretation, all sensitivity tests regard customer segment 1 only.

The results for Unsegmented MCCM versus Myopic ALC are displayed in Fig. 5, organized in a 3-by-3 matrix format. Across the columns, we display sensitivity to: cost of provision (left), ancillary relevance (center), and transition probability (right), while across the rows we display the corresponding impact on prices (top), purchase rates (center), and expected revenue (bottom).

- **Cost of provision:** The cost of provision for all three ancillary services remains equal and is increased from

$(c_1, c_2, c_3) = (\$0, \$0, \$0)$ to $(c_1, c_2, c_3) = (\$30, \$30, \$30)$. As costs increase, the offer prices increase non-linearly for both Myopic ALC and MCCM. Consequently, the purchase rates decrease. Notice that when the cost is low, MCCM strongly incentivizes customers to buy bundle offers by pricing the atomic ancillaries {1}, {2}, and {3} disproportionately high compared to Myopic ALC, which can be seen from the offer purchase rate distribution. Finally, MCCM systematically outperforms Myopic ALC, reaching a revenue gain of 23% under these ideal conditions. As cost increases, Myopic ALC and MCCM converge, leading to identical performance.

- **Relevance:** The relevance for all three ancillary services is increased from $(\phi_1, \phi_2, \phi_3) = (10\%, 10\%, 10\%)$ to $(\phi_1, \phi_2, \phi_3) = (100\%, 100\%, 100\%)$. The Myopic ALC price is indifferent to the relevance parameter since the relevance parameter only affects the volume of customers, not the price sensitivity of customers interested in an atomic ancillary. With MCCM pricing, relevance plays an important role, because the relevance for a bundled offer is the product of the atomic ancillaries' relevance. Hence when relevance is low, customers find bundles unattractive, causing prices to converge to the Myopic ALC prices; for high relevance, bundles become very attractive causing MCCM to price the atomic offers {1}, {2}, and {3} high and creating large discounts for the bundled offers. This trend is also reflected in the potential revenue benefit of MCCM, which increases strongly with passenger relevance.
- **Transition probability:** The transition probabilities remain equal $\rho_{ij} = \rho$ for $i \neq j, i, j = 1, \dots, N$ and are increased from $\rho = 0\%$ to $\rho = 16\%$. Thus, the transition probability from any offer O_i to the no-purchase terminal state $\rho_{i0} = 1 - 6\rho$, will decrease linearly from $\rho_{i0} = 100\%$ (immediate exit) to $\rho_{i0} = 4\%$. The case $\rho = 0\%$ essentially corresponds to the Myopic ALC case, given that all customers exit immediately after only considering one offer. This corner case recovers the Myopic ALC prices. At the other extreme when exit probability is very low, customers are captive and more likely to purchase an offer, which allows MCCM to capitalize on this fact and increase prices on all offers. Despite the price increase, MCCM expects increasing conversion rates, with the results being a sharp increase in expected revenue compared to the myopic ALC baseline.

Throughout the sensitivity test, as well as the other scenarios tested in this paper, we observe that the MCCM is not only scalable and customizable, but also that bundle prices remain logical and robust, which are desirable properties. Bundles were priced at a discount compared to the sum of their constituents, partially a result of our additive Normal valuation assumption. However, logical pricing is



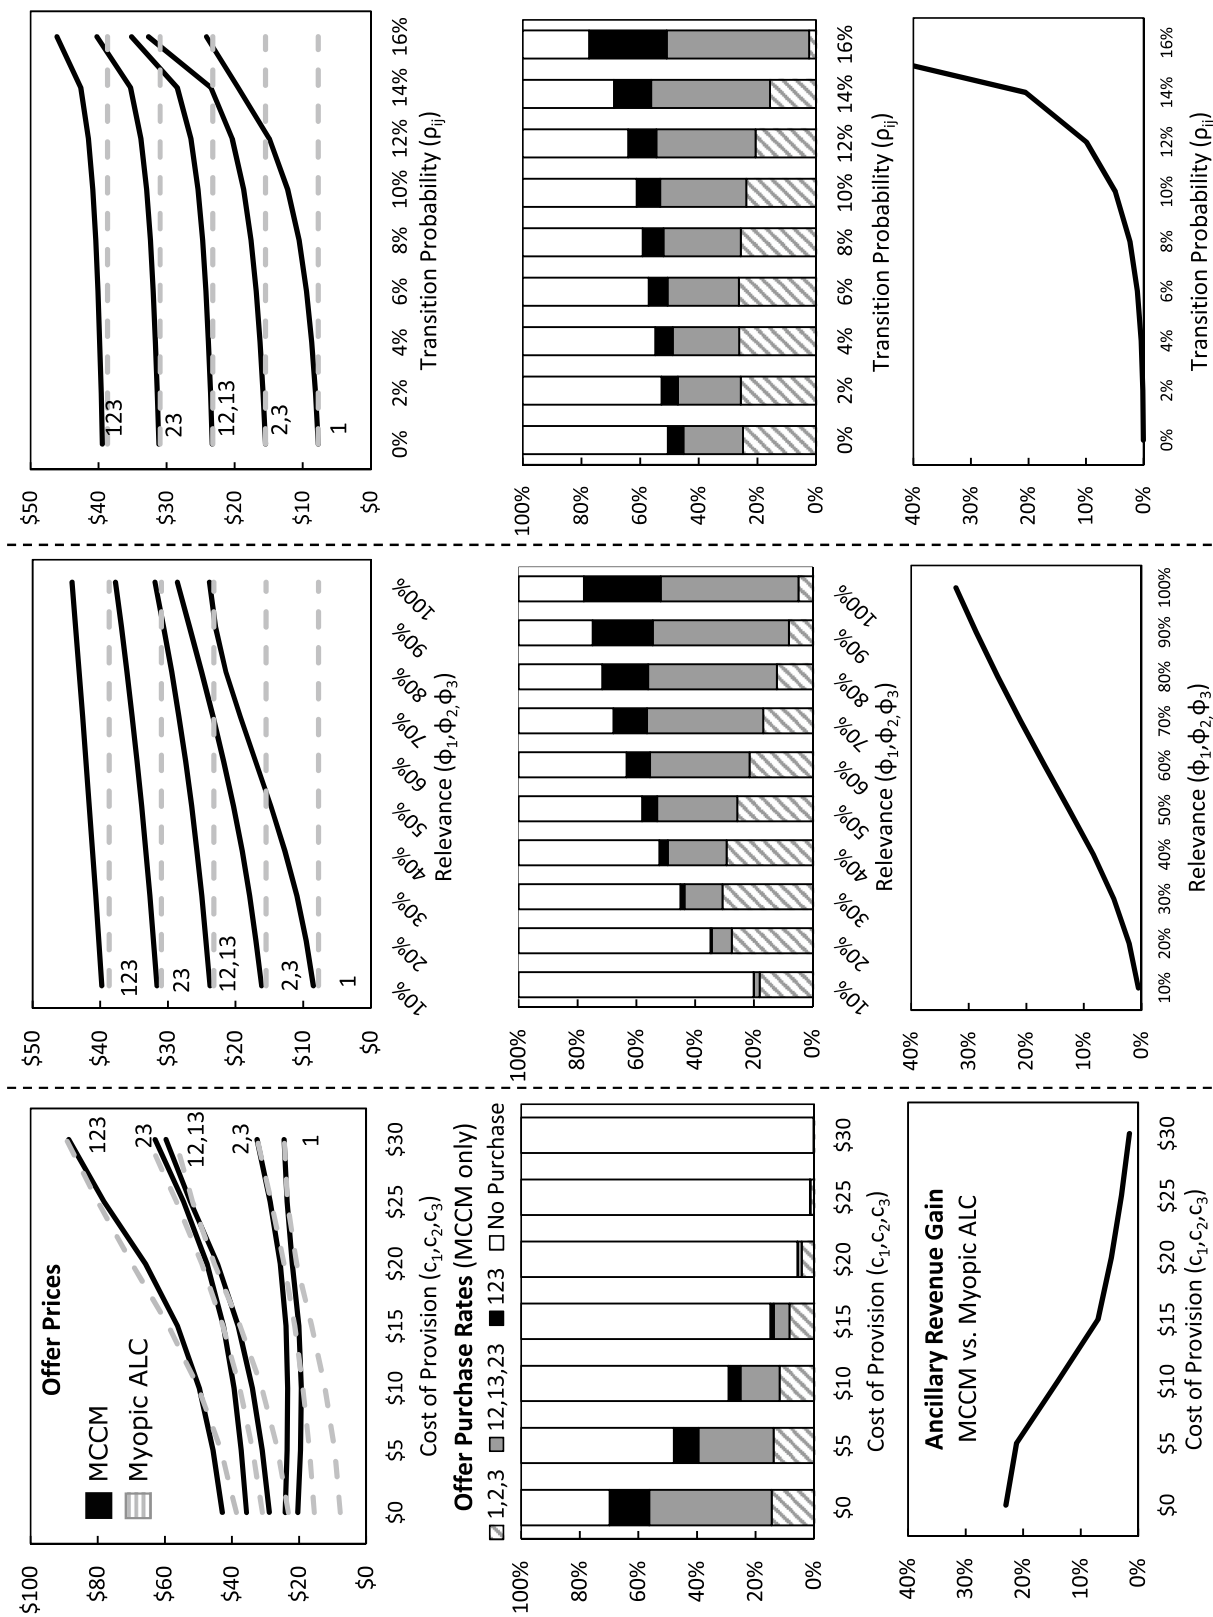


Fig. 5 Sensitivity of Markov chain choice (MCCM) and Myopic a la carte (Myopic ALC) prices to ancillary cost, relevance, and transition probability. Also shown is the expected purchase behavior under the MCCM prices and their relative contribution gains compared to the a la carte pricing baseline



not guaranteed for any combination of MCCM transition probabilities and valuations, as counterexamples can be constructed when products are valued at different orders of magnitude or when extreme values are used for transition probabilities.

Unsegmented offer set selection, Unsegmented pricing—(UU)

Table 5 expands the results from Table 4 by systematically evaluating all the 127 possible offer sets $S \subseteq \Omega$ using the unsegmented MCCM. For each offer set size $|S|$, we report the revenue maximizing offer set $S^* = \operatorname{argmax}_S \pi(S, \mathbf{p}(S))$, the corresponding offer prices $\mathbf{p}(S)$, and the expected purchase rates $\mathcal{P}(S)$ by customer segment. The full offer set $S = \Omega$, from Table 4, can be identified as the top row of Table 5.

Recall that as offers are removed from the full offer set, the customer choice behavior will adapt, with arrival and transition probabilities redistributed evenly across the remaining offers under the MCCM.

We first note that the optimal price for the same offer depends on the offer set. This is a major departure from both current industry practice and myopic ALC pricing, which results in a fixed price irrespective of offer set.

For example, in the offer set $S = \{\{2\}, \{2, 3\}, \{1, 2, 3\}\}$, the price of offer {2} is \$19.45, whereas the price for the same offer in the full offer set is \$21.71. This is because MCCM balances the contribution from a sale with (a) the risk of no purchase and (b) the potential buy-up to more profitable bundles. When only few alternatives exist, the risk of no purchase is increased, while if many alternatives exist the customer “stays longer in the shop” and the model takes a larger risk of proposing the more expensive (or consequently more profitable) offers.

Observe that according to Table 5, the offer set S^* with the highest expected contribution is of size 6 with $S^* = \Omega \setminus \{1\}$. This is because offer {1} is highly unattractive as can be seen in the low conversion rates. Eliminating unattractive offers reduces the risk of no-purchase.

Another use case relates to offer set selection for mobile devices such as phones and tablets. Their display size constrains the number of offers that can be displayed at a time. As an example, if we limit the offer set size to $|S| = 3$, we observe that the optimal offer set becomes $S = \{\{2\}, \{2, 3\}, \{1, 2, 3\}\}$, which is intuitive because all offers include the most relevant ancillary {2} for both customer segments.

Table 5 Best unsegmented offer set and MCCM prices for each offer set size from one to seven, tested with two customer segments with different offer relevance

$ S $		{1}	{2}	{3}	{1, 2}	{1, 3}	{2, 3}	{1, 2, 3}	$\mathcal{P}(S)$	$\pi(S)$
7	Price	\$17.00	\$21.71	\$21.77	\$27.39	\$27.40	\$34.54	\$42.11		\$16.71
	Seg 1	0.3%	10.4%	10.3%	13.2%	13.2%	15.1%	12.3%	74.9%	\$22.21
	Seg 2	0.0%	23.8%	6.0%	1.9%	0.5%	11.1%	0.6%	43.9%	\$11.21
6	Price	-	\$21.61	\$21.69	\$27.33	\$27.34	\$34.47	\$42.07		\$16.95
	Seg 1	-	10.7%	10.6%	13.5%	13.5%	15.4%	12.5%	76.2%	\$22.55
	Seg 2	-	24.4%	6.1%	1.9%	0.5%	11.2%	0.6%	44.7%	\$11.36
5	Price	-	\$21.19	-	\$27.07	\$27.08	\$34.24	\$41.90		\$16.78
	Seg 1	-	12.7%	-	15.3%	15.3%	17.3%	13.9%	74.4%	\$22.69
	Seg 2	-	27.6%	-	2.1%	0.5%	11.9%	0.6%	42.7%	\$10.88
4	Price	-	\$20.41	-	\$26.59	-	\$33.82	\$41.60		\$16.51
	Seg 1	-	16.4%	-	18.3%	-	20.3%	16.1%	71.1%	\$21.79
	Seg 2	-	31.2%	-	2.1%	-	12.0%	0.6%	45.9%	\$11.23
3	Price	-	\$19.45	-	-	-	\$33.29	\$41.22		\$16.04
	Seg 1	-	22.4%	-	-	-	24.9%	19.3%	66.6%	\$20.60
	Seg 2	-	36.7%	-	-	-	12.3%	0.6%	49.6%	\$11.48
2	Price	-	-	-	-	-	\$32.47	\$40.62		\$13.55
	Seg 1	-	-	-	-	-	33.1%	25.0%	58.1%	\$20.89
	Seg 2	-	-	-	-	-	18.0%	0.9%	18.9%	\$6.21
1	Price	-	-	-	-	-	\$31.16	-		\$10.61
	Seg 1	-	-	-	-	-	47.9%	-	47.9%	\$14.92
	Seg 2	-	-	-	-	-	20.2%	-	20.2%	\$6.30



Table 6 Segmented, optimized MCCM prices, and best offer sets for each customer segment and their expected revenue and offer purchase rates for each segment choosing under their MCCM

Segment 1 (FS & SS)									Segment 2 (FS & SS)																		
	$ S $	{1}	{2}	{3}	{1, 2}	{1, 3}	{2, 3}	{1, 2, 3}	$\pi(S), \mathcal{P}(S)$		$ S $	{1}	{2}	{3}	{1, 2}	{1, 3}	{2, 3}	{1, 2, 3}	$\pi(S), \mathcal{P}(S)$								
7	Price	\$20.41	\$24.18	\$24.18	\$28.89	\$28.89	\$35.66	\$42.88	\$22.45	7	\$12.43	\$18.67	\$18.88	\$25.63	\$25.64	\$33.11	\$41.09	\$11.73	7	0.6%	31.9%	7.8%	1.9%	0.5%	10.2%	0.5%	53.4%
	Prob.	0.0%	7.3%	7.3%	12.9%	12.9%	16.0%	13.5%	69.9%		0.6%	31.9%	7.8%	1.9%	0.5%	10.2%	0.5%	53.4%									
6	Price	-	\$24.10	\$24.10	\$28.81	\$28.81	\$35.58	\$42.81	\$22.79	6	-	\$18.59	\$18.83	\$25.60	\$25.61	\$33.07	\$41.07	\$11.87	6	-	7.4%	7.4%	13.1%	13.1%	16.3%	13.6%	70.9%
	Prob.	-	7.4%	7.4%	13.1%	13.1%	16.3%	13.6%	70.9%		-	32.9%	7.9%	1.9%	0.5%	10.4%	0.5%	54.1%									
5	Price	-	\$23.76	-	\$28.57	\$28.57	\$35.37	\$42.65	\$22.88	5	-	\$18.45	\$18.74	\$25.55	-	\$33.00	\$41.04	\$12.00	5	-	8.6%	-	14.5%	14.5%	17.7%	14.8%	70.1%
	Prob.	-	8.6%	-	14.5%	14.5%	17.7%	14.8%	70.1%		-	34.1%	8.1%	1.9%	-	10.5%	0.5%	55.1%									
4	Price	-	-	-	\$28.25	\$28.25	\$35.09	\$42.43	\$23.09	4	-	\$18.24	\$18.59	\$25.46	-	\$32.90	-	\$12.14	4	-	-	-	\$27.38	-	\$34.36	\$41.89	\$22.18
	Prob.	-	-	-	16.4%	16.4%	19.8%	16.3%	68.9%		-	36.0%	8.3%	2.0%	-	10.7%	-	57.0%									
3	Price	-	-	-	\$27.38	-	\$34.36	\$41.89	\$22.18	3	-	\$17.89	\$18.36	-	-	\$32.73	-	\$12.32	3	-	-	-	20.8%	-	24.2%	19.4%	64.4%
	Prob.	-	-	-	20.8%	-	24.2%	19.4%	64.4%		-	39.6%	8.8%	-	-	11.1%	-	59.5%									
2	Price	-	-	-	-	-	\$33.33	\$41.11	\$20.92	2	-	\$17.19	-	-	-	\$32.41	-	\$12.37	2	-	-	-	-	-	32.1%	24.9%	57.0%
	Prob.	-	-	-	-	-	32.1%	24.9%	57.0%		-	48.8%	-	-	-	12.3%	-	61.1%									
1	Price	-	-	-	-	-	\$31.16	-	\$14.92	1	-	\$15.48	-	-	-	-	-	\$11.39	1	-	-	-	-	-	47.9%	47.9%	
	Prob.	-	-	-	-	-	47.9%	-	47.9%		-	73.6%	-	-	-	-	-	73.6%									

Full offer set or segmented offer set selection, Segmented pricing—(FS and SS)

We conclude with segmented pricing tests. The simulation is now performed using two different MCCM models—one for each segment with the appropriate parameters, see Table 1. The results are displayed in Table 6 split by segment. The top row of Table 6 shows the results of segmented pricing with the full offer set (scenario FS). The optimal offer set with segmented offer set selection (scenario SS) is highlighted separately for each segment. For segment 1, the optimal offer $S_1^* = \{\{1, 2\}, \{1, 3\}, \{2, 3\}, \{1, 2, 3\}\}$ only contains bundled offers, while the optimal offer for segment 2 is quite different $S_2^* = \{\{2\}, \{2, 3\}\}$. The overall expected revenue per passenger across both segments is \$17.09 for scenario FS and \$17.73 for scenario SS, which is the highest among all the investigated scenarios.

First, we can clearly see the impact of segmented pricing within the full offer set (FS), where all offers differ in price between the two segments. While the price for the full bundle $\{1, 2, 3\}$ is relatively similar (\$42.88 for segment 1 and \$41.09 for segment 2), atomic offers such as $\{2\}$ are priced very differently (\$24.18 for segment 1 and \$18.67 for segment 2). Although the valuations for the offers are the same for both segments, they differ in relevance: offer $\{1, 2, 3\}$ is relevant to $\phi_1(\{1, 2, 3\}) = \phi_{1,1}\phi_{2,1}\phi_{3,1} = 42\%$ of segment 1 but only $\phi_2(\{1, 2, 3\}) = \phi_{1,2}\phi_{2,2}\phi_{3,2} = 1.2\%$ of segment 2. As a result, MCCM strongly encourages segment 1 to buy up to a bundle, whereas segment 2 is instead offered attractive

prices for the atomic offers. The impact of the different pricing strategies is clearly reflected in the expected purchase rates, with 7.3% of segment 1 and 31.9% of segment 2 expected to purchase $\{2\}$.

These results of segmented offer set selection are also quite logical. Since segment 1 finds bundle offers relevant, MCCM removes the less profitable atomic offers. Because segment 2 has almost no interest in ancillary a_1 , a strong interest in ancillary a_2 and moderate interest in ancillary a_3 , MCCM removes all offers containing a_1 and only allows bundling between a_2 and a_3 . These are important and desirable properties that allows us to use MCCM to tailor offer sets at the individual (segmented) customer level.

Notice that even within the segmented offer sets, the price for the same offer $\{2, 3\}$ differs (\$35.09 for segment 1 and \$32.41 for segment 2). Because segment 1 could further upsell to the offer $\{1, 2, 3\}$, the offer is priced higher than for segment 2, for whom it represents the most profitable offer and is, therefore, priced more attractively.

Summary of simulation results

Figure 6 compares the expected revenue performance π^* across the investigated scenarios. The figure displays the revenue per passenger in three groups, for the overall revenue (left hand side) and for the customer segments separately (right hand side). From the baseline (FM), we can assess the revenue increment of implementing gradually more sophistication into dynamic offer creation. Simply



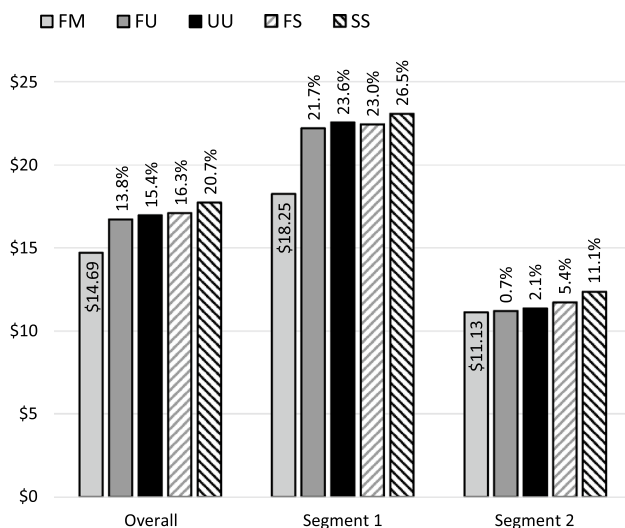


Fig. 6 Summary of expected revenue per passenger across the different scenarios, following the nomenclature used in Table 1. First letter represents Offer Set Selection (Full, Unsegmented, Segmented). Second letter represents Offer Set Pricing (Myopic, Unsegmented, Segmented)

introducing MCCM (scenario FU) and moving from the myopic approach to optimizing prices jointly for the entire offer set leads to a +13.8% increase in expected ancillary revenue under ideal conditions.

As we increase sophistication, allowing the model to perform unsegmented offer set selection (scenario UU), we obtain an additional increment of +1.6pp. Introducing segmented pricing with the full offer set (scenario FS) further increases revenue by +0.9pp. Finally introducing segmented offer set selection alongside segmented pricing (scenario SS) further increases revenue by +4.4pp to a total expected ancillary revenue uplift of +20.7%. The ancillary revenue uplift is dominated by segment 1, because segment 1 has higher affinity for bundled offers, whereas our baseline a la carte pricing is closer to optimality for segment 2.

Here, the revenue benefit of segmentation from UU to SS is relatively low, because in this study, we deliberately chose the same ancillary valuation across segments. Past research in the dynamic pricing of flight itineraries has shown that segmentation is particularly effective when there are large differences in valuation between segments (Wittman and Belobaba 2018), and we would expect even stronger performance from SS had we tested an environment with differentiated valuation parameters.

Note that in our results, based on the high ancillary relevance we chose, the overall ancillary purchase rate and revenue per passenger are likely higher than in reality. We also made simplifying assumptions: we assumed that customers

choose their offers according to an MCCM exactly, that removing unattractive offers from an offer set can change their purchase behavior and redistribute their transition probability across the remaining offers, and that an airline is able to perfectly differentiate between two customer segments. If the MCCM is not calibrated to accurately represent realistic customer behavior, then the revenue benefits could be diminished.

Conclusions

This paper investigated a tractable scientific model to dynamically create and price offers of airline ancillary services. We have demonstrated through a proof of concept that the Markov chain choice model (MCCM) could be used to solve this problem at scale.

While the MCCM has been applied to other assortment optimization and pricing problems before, we are the first to use it for bundle pricing and dynamic offer creation. We made several enhancements as we adapted the MCCM to this setting. We introduced a structured definitional and notational framework for offers and offer sets which had been absent in the academic literature. We also introduced the concept of ancillary relevance into the MCCM, which penalizes our model for showing offers that are not of interest to the customer.

We found that the MCCM possesses desirable properties of a discrete choice model for this problem setting—the price optimization is computationally tractable and produces intuitive prices for a variety of offer sets, including discounts for bundle offers. The model’s structure allows for efficient optimization of offer prices, even for larger offer sets. Its main limitation lies in the large number of parameters, especially for larger offer sets, which makes it challenging to estimate.

We showed that the MCCM can adjust its pricing, depending on the passenger valuations and relevance, to either favor a la carte or bundle pricing. Compared to an a la carte pricing model, the MCCM enables revenue gains from bundle pricing when the circumstances are favorable. Our results showed that the model can not only recommend different offer prices for different customer segments, but also entirely different offer sets, leading to revenue gains over unsegmented offer set selection and pricing.

Our work represents the first demonstration that the segmented dynamic offer creation problem can indeed be solvable in practice. We believe that in solving the offer set price optimization problem efficiently with the MCCM, we remove a major scientific roadblock for airlines to provide



a tailored, modern shopping experience to their customers, although further obstacles to its implementation remain.

Future work

This paper demonstrated that under ideal conditions, there can be significant revenue opportunities to be exploited through dynamic offer creation with our model. However, our revenue results assume that passengers will choose exactly according to the prescribed MCCM. Further simulation studies are required to understand how well the MCCM can represent realistic passenger choice behaviors and whether the revenue gains are sustained.

Furthermore, since the MCCM has significantly more parameters than many other discrete choice models, developing efficient and sufficiently accurate parameter estimation methods requires additional research. Further investigation is also needed to explore how the MCCM performs with estimated parameters, in terms of how well the model can predict customer choice behavior in reality, the offer sets and prices it recommends, and the overall revenue performance. Performance of the parameter estimation in situations where the available training data are limited is also an area for future inquiry, as airlines today often have limited information about the ancillary services that are purchased by their customers.

In addition, we have not investigated in detail the computational speed of performing offer set selection using our approach. This would be an important focus for a real-life airline implementation, as the optimal offer set would need to be returned nearly instantly in response to a search. Efficient search algorithms may be needed to speed up the offer set selection process in cases with very large offer sets. Finally, accurate customer segmentation models and guardrails to ensure customers receive the correct targeted offers would be needed to benefit from segmented offer set selection and pricing in a real-life airline application.

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, as well as Alexander Papen for verifying our results.

Funding Open Access funding provided by the MIT Libraries.

Data availability The data supporting this article are available from the corresponding author upon reasonable request.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated

otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

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.
- Alptekinoglu, A., and J.H. Semple. 2016. The exponential choice model: A new alternative for assortment and price optimization. *Operations Research* 64 (1): 79–93.
- Aydin, G., and J.K. Ryan. 2000. Product line selection and pricing under the multinomial logit model. *Proceedings of the 2000 MSOM Conference*.
- Belobaba, P. 1987. Air travel demand and airline seat inventory management. Massachusetts Institute of Technology, Ph.D. thesis. <https://dspace.mit.edu/bitstream/handle/1721.1/14800/17391833-MIT.pdf>.
- Berbeglia, G. 2016. Discrete choice models based on random walks. *Operations Research Letters* 44 (2): 234–237.
- Berbeglia, G., A. Garassino, and G. Vulcano. 2021. A comparative empirical study of discrete choice models in retail operations. *Management Science* 68: 4005–4023.
- Blackwell, D. 1965. Discounted dynamic programming. *Annals of Mathematical Statistics* 36: 226–235.
- Blanchet, J., G. Gallego, and V. Goyal. 2016. A Markov chain approximation to choice modeling. *Operations Research* 64 (4): 886–905.
- Calzada, J., and F. Martinez-Santos. 2016. Pricing strategies and competition in the mobile broadband market. *Journal of Regulatory Economics* 50: 70–98.
- CAPA, 2018. Airline outlook: Unbundling, fare families, ancillaries. Centre for Asia Pacific Aviation. <https://centreforaviation.com/analysis/reports/airline-outlook-unbundling-fare-families-ancillaries-447478>.
- Chu, C.S., P. Leslie, and A. Sorensen. 2008. *Nearly Optimal Pricing for Multiproduct Firms*. National Bureau of Economic Research Working Paper 13916.
- Chu, C.S., P. Leslie, and A. Sorensen. 2011. Bundle-size pricing as an approximation to mixed bundling. *American Economic Review* 101 (1): 263–303.
- Dadoun, A., M. Defoin-Platel, T. Fiig, C. Landra, and R. Troncy. 2021. How recommender systems can transform airline offer construction and retailing. *Journal of Revenue and Pricing Management* 20: 301–315.
- Daft, J., S. Albers, and S. Stabenow. 2021. From product-oriented flight providers to customer-centric retailers: A dynamic offering framework and implementation guidelines for airlines. *Journal of Revenue and Pricing Management* 20: 615–625.
- Dong, J., A.S. Simsek, and H. Topaloglu. 2019. Pricing problems under the Markov chain choice model. *Production and Operations Management* 28 (1): 157–175.
- Elberse, A. 2010. Bye-bye bundles: The unbundling of music in digital channels. *Journal of Marketing* 74: 107–123.
- Ettl, M., P. Harsha, A. Papush, and G. Perakis. 2020. A data-driven approach to personalized bundle pricing and recommendation. *Manufacturing & Service Operations Management* 22 (3): 461–480.
- 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., K. Isler, C. Hopperstad, and P. Belobaba. 2010. Optimization of mixed fare structures: Theory and applications. *Journal of Revenue and Pricing Management* 9: 152–170.
- Fiig, T., K. Isler, C. Hopperstad, and S.S. Olsen. 2012. Forecasting and optimization of fare families. *Journal of Revenue and Pricing Management* 11: 322–342.
- Fiig, T., R. Le Guen, and M. Gauchet. 2018. Dynamic pricing of airline offers. *Journal of Revenue and Pricing Management* 17: 381–393.
- Fuerderer, R., A. Huchzermeier, and L. Schrage. 1999. Stochastic option bundling and bundle pricing. In *Optimal bundling*, ed. R. Fuerderer, A. Herrmann, and G. Wuebker, 61–86. Berlin, Heidelberg: Springer.
- Gallego, G., and C. Stefanescu. 2010. Service Engineering: The Future of Service Design and Pricing. Columbia University Center for Pricing and Revenue Management Working Paper Series No. 2010–02.
- Gallego, G., and R. Wang. 2014. Multiproduct price optimization and competition under the nested logit model with product-differentiated price sensitivities. *Operations Research* 62 (2): 450–461.
- Gupta, A., and D. Hsu. 2020. Parameter identification in Markov chain choice models. *Theoretical Computer Science* 808: 99–107.
- Hanson, W., and R.K. Martin. 1990. Optimal bundle pricing. *Management Science* 36 (2): 155–174.
- Hjorth, R., T. Fiig, J. Larsen, and N. Bondoux. 2018. Joint overbooking and seat allocation for fare families. *Journal of Revenue and Pricing Management* 17: 436–452.
- IATA. 2021. Airline retailing: An industry vision for Offers and Orders. IATA Air Retailing White Paper. <https://www.iata.org/contentassets/47c4d32973014560beef3cc421bdf402/airline-retailing-an-industry-vision-for-offers-and-orders.pdf>.
- Kobayashi, B.H. 2005. Does economics provide a reliable guide to regulating commodity bundling by firms? A survey of the economic literature. *Journal of Competition Law & Economics* 1 (4): 707–746.
- Kumar, R., A. Li, and W. Wang. 2018. Learning and optimizing through dynamic pricing. *Journal of Revenue and Pricing Management* 17: 63–77.
- Kumara, M.S.R., B.R. Guntreddy, I.G. Vega, and Y.H. Tai. 2021. Dynamic pricing of ancillaries using machine learning: One step closer to full offer optimization. *Journal of Revenue and Pricing Management* 20: 646–653.
- Loughry, J., J.I. van Hemert, and L. Schoofs. 2000. Efficiently enumerating the subsets of a set. <http://www.applied-math.org/subset.pdf>.
- Ma, W. 2021. When is assortment optimization optimal? <https://arxiv.org/abs/2105.10433>.
- Madireddy, M., R. Sundararajan, G. Doreswamy, M.H. Nia, and A. Mital. 2017. Constructing bundled offers for airline consumers. *Journal of Revenue and Pricing Management* 16: 532–552.
- Ødegaard, F., and J.G. Wilson. 2016. Dynamic pricing of primary products and ancillary services. *European Journal of Operations Research* 251: 586–599.
- Ratliff, R., and G. Gallego. 2013. Estimating sales and profitability impacts of airline branded-fares product design and pricing decisions using customer choice models. *Journal of Revenue and Pricing Management* 12: 509–523.
- Schmalensee, R. 1984. Gaussian demand and commodity bundling. *Journal of Business* 57: S211–S230.
- Schubert, D., C. Sys, and R. Macário. 2021. Customized airline offer management: A conceptual architecture. *Journal of Revenue and Pricing Management* 21: 553–563.
- 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.
- Simsek, A.S., and H. Topaloglu. 2018. Technical note: An expectation-maximization algorithm to estimate the parameters of the Markov chain choice model. *Operations Research* 66 (3): 748–760.
- Sorensen, J. 2021. The 2021 CarTrawler yearbook of ancillary revenue. IdeaWorksCompany Airline Revenue and Transformation Series.
- Talluri, K., and G. van Ryzin. 2004. Revenue management under a general discrete choice model of customer behavior. *Management Science* 50 (1): 15–33.
- Touraine, S., and H. Coles. 2018. Dynamic Offer Creation. IATA Air Industry Retailing White Paper. <https://www.iata.org/contentassets/0688c780d9ad4a4fad461b479d64e0d/air-white-paper-dynamic-offer-creation.pdf>.
- Touraine, S. 2021. The industry transformation to dynamic offering. *Journal of Revenue and Pricing Management* 20: 611–614.
- 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.
- Walczak, D., and E. Kambour. 2014. Revenue management for fare families with price-sensitive demand. *Journal of Revenue and Pricing Management* 3: 273–290.
- Wang, K. 2021. Ancillary Price Optimization with Thompson Sampling - Initial Tests. *Presented at MIT PODS Consortium Conference*, November 2021.
- Wang, K. 2022. Total revenue optimization approaches with the Markov chain choice model. *Proceedings of the 2022 AGIFORS RM Study Group Conference*.
- Wang, K.K., M.D. Wittman, and A. Bockelie. 2021. Dynamic offer generation in airline revenue management. *Journal of Revenue and Pricing Management* 20: 654–668.
- Wittman, M.D., and P.P. Belobaba. 2018. Customized dynamic pricing of airline fare products. *Journal of Revenue and Pricing Management* 17: 78–90.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Kevin K. Wang 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.

Michael D. Wittman 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.

Thomas Fiig is the Director, Chief Scientist at Amadeus. He holds a Ph.D. in Theoretical Physics and Mathematics from the University of Copenhagen, Denmark. He has more than 20 years of revenue management experience and has published several papers on forecasting, optimization of simplified fare structures, dynamic pricing, offer optimization, and recently in various areas of machine learning.

