Flexible Services and Manufacturing Journal
https://doi.org/10.1007/s10696-024-09545-x

On the value of booking data for upsell decision-making in revenue management

Davina Rauhaus[1] · Jochen Gönsch[1] · Claudius Steinhardt[2]

Accepted: 15 May 2024

Abstract
In passenger aviation and many other areas of transportation, it is common practice to offer customers who have booked a ticket for a lower compartment free seats in higher compartments at a discount before departure, a practice known as upselling. For example, economy class customers are offered a seat in business class for a small surcharge a few days before take-off. Obviously, it matters to whom to offer an upsell and at what price. In this paper, we address this decision problem in a generic fashion for revenue management settings. We assume that the company has disaggregated booking data about the customer's initial choice of a product from a provided offer set. This data contains information about individual customers' preferences and may be leveraged to decide on upsell prices. To this end, we propose an optimization approach based on an expectation model, in which customers' response probability is represented as a conditional probability formally consistent with their initial buying decision in a multinomial logit model. We present variants of the approach based on different levels of exploitable customer-specific booking data. In a numerical study, we investigate the value of this data usage and upselling in general to the company. Upselling in conjunction with knowledge of the customers' original offer sets and customer segments, substantially increases revenues. Furthermore, the study demonstrates that the proposed approach can lead to larger revenue benefits than a naive benchmark approach which statistically decouples the customers' upgrade acceptance decision from their original choice during the purchasing process.

Keywords Upsell · Upgrade · Pricing · Revenue management · Customer choice · Multinomial logit

Extended author information available on the last page of the article

1 Introduction

Many service companies like airlines, hotels, and railways dispose of vertically differentiated resources with a fixed capacity (e.g., seats in economy, business, and first-class compartments for airlines). If mismatches between supply and demand occur, capacity in a higher-valued compartment may be used to substitute missing capacity in a lower-valued one (e.g., a passenger with an economy ticket obtains a seat in business class).

An easy way to provide this substitution is using upgrades, where the customer (she/they) receives the better product at no extra charge and, thus, is assumed to always be happy to accept it. Obviously, this additional flexibility can be best leveraged if it is adequately modeled in the revenue management systems that dynamically decide on the offer set—i.e., the assortment and pricing of products—presented to incoming customer requests during the booking horizon. The integration of upgrades into revenue management's capacity control has been intensively researched (see Gönsch and Steinhardt 2015 for a review).

About a decade ago, companies started to monetize on upgrades. They discovered that customers who initially bought a lower-valued product may be willing to pay for the better one. Thus, companies may offer them an upsell at a later point in time, enabling the customer to switch to the higher-valued product for a small surcharge (see Fig. 1 for an example). Basically, an upsell means that a customer is offered an upgrade at a certain fee. So, the question arises to whom and at what price an upsell should be offered. Although upsells are widely used today, this decision is made

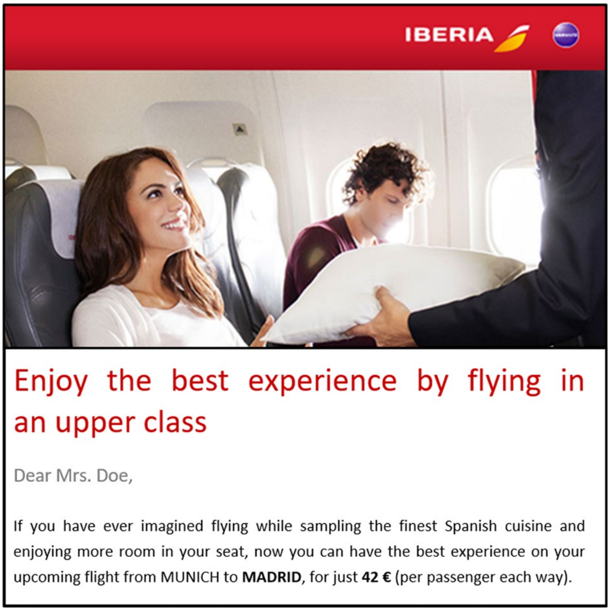

Fig. 1 Exemplarily e-mail with upsell offer from IBERIA

rather hands-on in practice and has—to the best of our knowledge—received very little attention in the scientific literature yet. One neglected aspect in the literature is the detailed consideration of customer-specific sales data. However, there is substantial potential in using such data, as the company knows more about the customer at this later stage, e.g., can derive some of their preferences from their choice from the original offer set. For this purpose, nowadays, due to the low prices of storage capacities, companies usually even have all the raw sales data available.

With this paper, we address this gap by investigating how companies can leverage booking data to improve upsell pricing decisions. Particularly, our contributions are as follows:

- First, we derive a closed-form representation of the customers' response probability to upsell offers consistent with their original buying decision from a multinomial logit model (MNL).
- Based upon this, we then propose an optimization model, more precisely, an expectation model to derive optimal upsell pricing decisions. The model is able to consider different levels of detail of original booking data.
- We also develop a benchmark approach that does not learn from the booking horizon at all, i.e., it treats the upsell response as statistically decoupled from the original purchasing choice. We prove that the resulting upsell probabilities systematically overestimate the actual ones.
- Finally, in a comprehensive numerical study, we are able to show that using more detailed, customer-specific booking data in the upsell optimization problem can increase revenues significantly, and that the benchmark approach performs always worse than using detailed data.

The remaining paper is structured as follows. Section 2 presents the relevant literature. In Sect. 3, the mathematical upsell model is developed, including the derivation of the upsell response probability. Further, for benchmarking purposes, we present a naïve approach that neglects any customer data from the original booking. Finally, we discuss the different levels of detail in the booking data that can be fed into the model, which we refer to as different states of information (SOI). Section 4 is dedicated to the numerical study. The paper closes with managerial insights and a conclusion as well as an outlook for further research in Sect. 5. Note that the research presented is directly applicable to many service industries where revenue management is applicable and upsells can be offered.

2 Literature overview

In the following, we first present survey papers and books on the related fundamental topics, more precisely, on revenue management, the MNL model and upgrades (Sect. 2.1). Next, we consider papers that treat upsells from a perspective other than that of revenue management (2.2), before we delineate upselling papers in revenue management from our paper (2.3). Finally, we refer to revenue management papers that use the term "upsell" differently (Sect. 2.4) to clarify the terminology.

2.1 Fundamental topics

Extensive literature exists on revenue management in general. For broad reviews of the field, see for example Belobaba (1987), Weatherford and Bodily (1992), and Chiang et al. (2007). In addition, some reviews concentrate on specific topics, e.g., generalizations and industry applications (Klein et al. 2020), price models (Bitran and Caldentey 2003), demand dependency (Weatherford and Ratliff 2010), and choice-based revenue management (Strauss et al. 2018). This paper concentrates on upsells and customer choice behavior based on a latent class model. Therefore, we limit the following literature review to these two topics.

Talluri and Van Ryzin (2004a) were the first to apply the multinomial logit model (MNL) to revenue management. In the econometric literature, the MNL model had been analyzed extensively by McFadden (1987). Using more than one customer segment leads to a so-called latent class model, like Bront et al. (2009) use in their approach for choice-based network revenue management. We refer to Strauss et al. (2018) for an overview of papers that deploy these discrete customer choice methods in revenue management.

Our topic is strongly related to upgrades. At this point, we omit a literature survey on (free) upgrades and refer to Gönsch and Steinhardt (2015) who review the literature regarding upgrades extensively. Some papers also discuss pricing with considering upgrades in the pricing process, e.g., Li et al. (2016, 2022). Upsells, in contrast to upgrades, are a relatively unexamined topic in revenue management.

2.2 Upsells in different disciplines

The following streams of literature add to the understanding of customer behavior concerning upsells, but do not provide companies with pricing strategies for upselling.

Ahn et al. (2022) consider upsells from a mainly econometric perspective. The authors compare emotional to rational buying decisions in online and offline upselling processes. Similarly, Denizci Guillet et al. (2022) investigate whether firms should use so-called online or offline upselling or even both strategies at the same time. Norvell et al. (2018) study the short-term and long-term impacts of upselling and down-selling from a field study. They find that, in retail, down-selling leads to customer loyalty and, on the other hand, upselling increases short-term revenues, but does not increase customer loyalty. Park and Yoon (2022) add to this literature by considering multi-brand retailers and examining the effects of promotions for high-end and lower-end brands. The last two mentioned papers include selling goods instead of services and investigate the effects of promotion.

Another perspective is given by Mayer et al. (2022): their paper examines the effect of option framing and cognitive load on customer choices in the tourism industry. Hence, their research can be classified as psychological and economic aspects of upselling.

2.3 Upsells in revenue management

Gallego and Stefanescu (2009) present a dynamic programming formulation for capacity control with planned upgrades. They focus on static deterministic linear programming (DLP) approximations to solve the model; one for a primary capacity provider and another one for a reseller. Moreover, they include customer choice models by defining corresponding upsell response probabilities and show that it is then no longer beneficial to offer upsells when capacity mismatches are known from the beginning of the booking horizon, when reseller margins are homogeneous, or when companies have full flexibility in pricing.

Çakanyıldırım et al. (2020) implement upsells into a dynamic program. They consider dynamic upsell offers from a regular to a premium product between the booking time and the check-in time. Their model contains decisions on price, time, and number of upsell offers. In addition, they model the firm's revenue maximization problem as a dynamic program and show that the optimal upgrade policy is of a pulsing type. They adhere to the independent demand assumption and use a first-come first-served policy during the booking horizon, i.e., no capacity control. This corresponds to an offer set always containing all available products and no upselling possibilities when the premium product is not offered anymore. In this case, it is not possible to exploit information from the booking horizon. Thus, they do not include this possibility in the optimization model.

Cui et al. (2019) examine the price dispersion when upsells are introduced at a price lower than the original price. The authors take an economic rather than a revenue management perspective. Besides, they only consider upgrades within the economy compartment (e.g., ancillary services). This differs from our approach in the way that we include capacity shifts between different compartments. Their main finding is that introducing upgrades supports price discrimination and results in a higher revenue for the company as well as in a higher consumer surplus.

Thirumuruganathan et al. (2023) focus on forecasting the upsell response probability using a machine learning algorithm based on real airline booking data. They provide insights into how customers can be segmented for upselling: those who never accept an upsell, those who always accept when offered, and those who consider taking the offer, but may not accept it because the price is too high. Although the authors use an integer linear program to assign upsells to customers, the model's results are price ranges with a certain likelihood that a customer accepts an upsell offer. Thus, they do not provide a specific price for offering an upsell. Additionally, compared to our approach, the authors do not consider the offer set for predicting the upsell response probability.

A related topic to upsells is conditional upgrades or standby upgrades. Although the customer pays for a conditional upgrade as for an upsell, she does not necessarily obtain the better product. These upgrades are sold in advance, but will be fulfilled if and only if capacity in higher compartments is not fully utilized at the end of the booking horizon. Exemplary papers considering conditional upgrades are Yilmaz et al. (2017, 2022), Cui et al. (2018) and Biyalogorsky et al. (2005).

2.4 Papers using the term upsells with a different meaning

The term upsell is not consistently used in literature. We define upsells as upgrades for which customers pay an extra fee. Nevertheless, we next discuss some papers that use the term differently.

Denizci Guillet (2020) investigates upsell auctions. Although this is related to our research, one major difference can be identified in the pricing mechanism: Instead of posted prices, the author prices the upsells by auctions, which leads to research to a game-theoretical problem.

Aydin and Ziya (2008), Shao (2021) as well as Bergemann et al. (2022) and Schaefers et al. (2022) use the term upselling as a kind of cross-selling. In their terminology, upselling takes place immediately after the purchase, and the term describes selling additional products instead of convincing the customer to switch to a better product.

Yet another definition is used by Gallego et al. (2009) and Seng Pun et al. (2016) who use the terms upsell and sell-up, respectively, to describe buy-ups, the opposite of a buy-down. A buy-up occurs if a customer who wants to buy a cheaper product buys a more expensive one because the cheaper one is unavailable.

3 Upsell price optimization

In the setting we consider, that the time horizon is partitioned into the selling period, the upsell horizon, and the service period (see Fig. 2), which do not overlap.

The sequence of events is as follows: The company first performs classical choice-based revenue management during the *selling period*. That is, customers sequentially arrive and may buy a product from the offer set presented to them. At the end of this period, the company often faces a capacity mismatch (e.g., free business class seats) that is unforeseen in the sense that revenue management could not anticipate or avoid it. Furthermore, for each customer who bought a product, the company now has precise information on bought products and maybe other information on the customers from the purchasing process, e.g., the customer segment and/or the seen offer set.

Then, the *upsell horizon* follows. This horizon is limited in time and short compared to the length of the selling period, e.g., to the last hours before the service is carried out. Here, the customer information can be used to efficiently sell the remaining capacity to the existing customers via upsells. To do so, at the beginning of the upsell horizon, an *upsell optimization problem* is solved using all available information to determine whom and at what price to offer upsells. Then, the company batch-wise offers upsells to the customers who individually decide whether to accept the upsell or not.

Finally, the *service period* begins, and the service is provided.

In this paper, we focus on the upsell horizon, in particular, on the upsell optimization model and on different possibilities to incorporate information on the customers obtained in the selling period.

In the following subsections, we first describe the problem formally and introduce the notation (Sect. 3.1). Section 3.2 presents the formal derivation of the upsell response probability. The expectation model for the upsell horizon is developed in Sect. 3.3. Subsequently, for benchmarking purposes, an alternative, naive approach is presented that does not learn from the booking horizon, but considers the customer's new choice situation as statistically decoupled from the original choice (Sect. 3.4). In Sect. 3.5, we describe the so-called states of information (SOI), that is, the level of detail used to record booking data. The section concludes with an extension of the upsell response probability to consider group bookings (Sect. 3.6).

3.1 Setting and notation

We build on a classical choice-based single-leg setting, distinguishing between multiple compartments. More precisely, the selling period is discretized into T micro periods, during which the company sells products from the set $\mathcal{J}$. To simplify notation, this set $\mathcal{J}$ includes the no-purchase alternative (indexed 0). Each product $j \in \mathcal{J}$ requires $a_{jh} \in \{0, 1\}$ units of capacity from resource (compartment) $h \in \mathcal{H} = \{1, 2, \dots, |\mathcal{H}|\}$. Lower (resp. higher) numbers refer to lower-valued (resp. higher-valued) resources. The (remaining) capacity of resource h is denoted by c_h. The price of product j is r_j. The no-purchase alternative does not lead to sales and, thus, has a price of $r_0 = 0$.

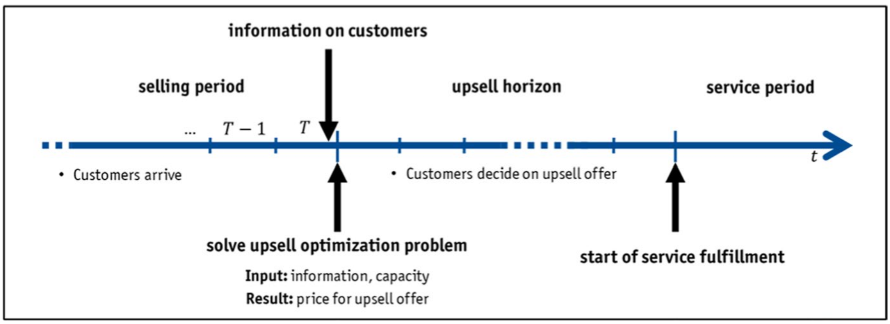

Fig. 2 Model timeline

In each micro period t, at most one customer arrives. In more detail, with probability λ_m, a customer belonging to customer segment $m \in \mathcal{M}$ arrives (w.l.o.g., we assume a time-homogeneous arrival process to improve readability). Without knowing the segment, the company offers her an offer set $\mathcal{O} \subseteq \mathcal{J}$ that takes into account the remaining capacity c_h at that time. The offer set always includes the no-purchase alternative $j = 0$, as the customer may decide not to buy. The customer chooses according to a random utility model, more precisely, according to the MNL model (see, e.g., McFadden 1987), which is based on random utility theory and has widely been used in revenue management (see, e.g., Talluri and Van Ryzin 2004a, b; Zhang and Adelman 2009; Rusmevichientong and Topaloglu 2012). This modeling framework requires the specification of utility functions. We choose a generic specification where a customer from segment m assigns a deterministic utility of $V_{jm} = \frac{1}{\beta_m} \cdot (q_{jm} - r_j)$ to product j. Here, q_{jm} is a segment-specific quality index that describes the preference of customer segment m for product j and β_m is a scale parameter determining the strength of the stochastic influence described below.

The econometric foundations of random utility theory imply that customers within one segment may differ according to individual tastes that are not observable. More precisely, a customer n from segment m chooses according to her stochastic (or total) utility $V_{jm} + \epsilon_{jn}$, with ϵ_{jn} being the realization of an i.i.d. Gumbel-distributed random variable in the MNL. If the customer decides against buying anything, she obtains the stochastic utility ϵ_{0n} of the no-purchase alternative which is always available. Now, if the company offers a customer n from segment m offer set $\mathcal{O} \subseteq \mathcal{J}$, she chooses a product j' with the highest stochastic utility. That is, a product j' with $V_{j'm} + \epsilon_{j'n} \geq V_{jm} + \epsilon_{jn} \forall j \in \mathcal{O}$ is chosen. If $V_{0m} + \epsilon_{0n}$ is the highest utility, no product is bought. For the MNL, it is well-known that the resulting probability $p_j(\mathcal{O})$ that she chooses product $j \in \mathcal{O}$ is given in closed form by the equation

$$p_j(\mathcal{O}) = \frac{e^{\frac{1}{\beta_m}(q_{jm} - r_j)}}{\sum_{i \in \mathcal{O}} e^{\frac{1}{\beta_m}(q_{im} - r_i)}} \quad \forall j \in \mathcal{O} \tag{1}$$

The quality indices q_{jm} as well as the scale parameter β_m are shared knowledge of customers and company, whereas her realizations of the ϵ are private knowledge of each customer.

In the subsequent upsell horizon, as usual in practice, the company offers at most one upsell to each customer. For that purpose, the customers $n \in \mathcal{N}$ who have bought a product are subclassified into classes $l \in \mathcal{L}$ such that all y_l customers in class l are homogeneous to the company, given the customer's data the company plans to use when making upsell pricing decisions. Since the company knows in each case which product the customers have bought, the classes are formed in such a way that all members of a class l have bought the same product j_l. The company may also know and exploit which segment customers belong to or which offer set they have seen, which means that in this case, classes would be formed with unique class-specific segments m_l or offer sets $\mathcal{O}_l$, respectively (see Sect. 3.5 for details).

As they are homogeneous to the company, we assume that all customers in a class either receive no upsell offer or the same. In the latter case, each customer in class l belonging to customer segment m_l receives exactly one upsell offer from product j_l to product k_l at price r_l. The conditional upsell response probability p_l is the probability that a customer from class l accepts the upsell offer. The upsell price is required not to exceed the difference between the two products‘ original prices, i.e., $0 \leq r_l < r_{k_l} - r_{j_l}$. If capacity is exceeded in the upsell horizon, the company can reject upsell requests.

3.2 Upsell response probability

Regarding the formulation of the upsell optimization problem, the probability that a customer buys an upsell is crucial. The following proposition provides this upsell response probability in closed form, conditional on the customer's behavior in the selling period. It is based on the assumption that the customer's unobservable preferences—represented by the realizations ϵ—do not change over time. Examples of non-observable, but constant preferences include, e.g., the height of the passengers or the weight of their luggage. Note that this is a generalization of the upsell response probability from Gallego and Stefanescu (2009). Moreover, we provide an explicit derivation.

Proposition 1 *The probability that a customer n from class l who initially bought product j_l out of offer set $\mathcal{O}_l$ accepts an upsell to k_l at price r_l is given by*

$$p_l = 1 - \frac{\sum_{i \in \mathcal{O}_l} e^{\frac{1}{\beta_{m_l}}(q_{im_l} - r_i)}}{\sum_{i \in \mathcal{O}_l \setminus \{k_l\}} e^{\frac{1}{\beta_{m_l}}(q_{im_l} - r_i)} + e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l} - (r_{j_l} + r_l))}} \qquad \forall \mathcal{O}_l \subseteq \mathcal{J}, l \in \mathcal{L} \tag{2}$$

Proof As a customer class's probability is independent of all other classes $l \in \mathcal{L}$, we ease notation in the proof by leaving out the index l. That is, we write β, O, q_i, j, and k instead of $\beta_{m_l}, \mathcal{O}_l, q_{im_l}, j_l$, and k_l. Furthermore, we introduce the abbreviation $w_{in} = q_i + \beta \cdot \epsilon_{in} \forall i \in \mathcal{J}, n \in \mathcal{N}$.

The starting point is rewriting the upsell response probability p_l as the complementary probability given in (3). Here, the second term is the conditional probability that a customer who bought product j out of offer set $\mathcal{O}$ now does not buy the upsell at price r_l.

$$p_l = 1 - \frac{P(\text{no upsell to } k \text{ at price } r_l \cap \text{ buying } j \text{ while seeing } \mathcal{O})}{P(\text{buying } j \text{ while seeing } \mathcal{O})} \tag{3}$$

The denominator is the probability that product j was bought by a customer from class l during the booking horizon given by (1).

In the numerator, we distinguish two cases: first, the upsell product k was part of the offer set ($k \in \mathcal{O}$), and second, it was not ($k \notin \mathcal{O}$). For $k \in \mathcal{O}$ the following holds:

$$\begin{aligned} &P(\text{no upsell to } k \text{ at price } r_l \cap \text{buying } j \text{ while seeing } \mathcal{O}) \\ &\quad = P\big(w_{jn} - r_j > w_{kn} - (r_j + r_l) \wedge w_{jn} - r_j > w_{in} - r_i \ \forall i \in \mathcal{O}\setminus\{j\}\big) \\ &\quad = P\big(w_{jn} - r_j > w_{kn} - (r_j + r_l) \wedge w_{jn} - r_j > w_{kn} - r_k \wedge w_{jn} - r_j > w_{in} - r_i \ \forall i \in \mathcal{O}\setminus\{j,k\}\big) \end{aligned}$$

Combining the first two inequalities yields:

$$\begin{aligned} &P(\text{no upsell to } k \text{ at price } r_l \cap \text{buying } j \text{ while seeing } \mathcal{O}) \\ &\quad = P\big(w_{jn} - r_j > w_{kn} - \min\{r_j + r_l, r_k\} \wedge w_{jn} - r_j > w_{in} - r_i \ \forall i \in \mathcal{O}\setminus\{j,k\}\big) \\ &\quad = \frac{e^{\frac{1}{\beta}(q_j - r_j)}}{e^{\frac{1}{\beta}(q_k - \min\{r_j + r_l, r_k\})} + \sum_{i \in \mathcal{O}\setminus\{k\}} e^{\frac{1}{\beta}(q_i - r_i)}} \end{aligned} \tag{4}$$

The last equality in (4) results from the fact that the consideration of the complementary probability allows us to apply the standard MNL closed-form transformation that can, e.g., be found in Ben-Akiva and Lerman (1985).

Finally, we substitute (1) and (4) in (3) and derive (5) for the first case ($k \in \mathcal{O}$):

$$\begin{aligned} p_l &= 1 - \frac{P(\text{no upsell to } k \text{ at price } r_l \cap \text{ buying } j \text{ while seeing } \mathcal{O})}{P(\text{buying } j \text{ while seeing } \mathcal{O})} \\ &= 1 - \frac{e^{\frac{1}{\beta}(q_j - r_j)}}{e^{\frac{1}{\beta}(q_k - \min\{r_j + r_l, r_k\})} + \sum_{i \in \mathcal{O}\setminus\{k\}} e^{\frac{1}{\beta}(q_i - r_i)}} \cdot \frac{\sum_{i \in \mathcal{O}} e^{\frac{1}{\beta}(q_i - r_i)}}{e^{\frac{1}{\beta}(q_j - r_j)}} \\ &= 1 - \frac{\sum_{i \in \mathcal{O}} e^{\frac{1}{\beta}(q_i - r_i)}}{e^{\frac{1}{\beta}(q_k - \min\{r_j + r_l, r_k\})} + \sum_{i \in \mathcal{O}\setminus\{k\}} e^{\frac{1}{\beta}(q_i - r_i)}} \end{aligned} \tag{5}$$

The second case covers $k \notin \mathcal{O}$. Now, the numerator does not include the minimum, because during the booking horizon, the customer did not compare products j and k. Thus, the event does not encompass $w_{jn} - r_j > w_{kn} - r_k$. Again, see Ben-Akiva and Lerman (1985) for the last step's transformation.

$$\begin{aligned} &P(\text{no upsell to } k \text{ at price } r_l \cap \text{buying } j \text{ while seeing } \mathcal{O}) \\ &\quad = P\big(w_{jn} - r_j > w_{kn} - (r_j + r_l) \wedge w_{jn} - r_j > w_{in} - r_i \ \forall i \in \mathcal{O}\setminus\{j\}\big) \\ &\quad = \frac{e^{\frac{1}{\beta}(q_j - r_j)}}{e^{\frac{1}{\beta}(q_k - (r_j + r_l))} + \sum_{i \in \mathcal{O}} e^{\frac{1}{\beta}(q_i - r_i)}} \end{aligned} \tag{6}$$

Substituting (1) and (6), formula (3) turns into (7):

$$
\begin{aligned}
p_l &= 1 - \frac{P(\text{no upsell to } k \text{ at price } r_l \cap \text{ buying } j \text{ while seeing } \mathcal{O})}{P(\text{buying } j \text{ while seeing } \mathcal{O})} \\
&= 1 - \frac{e^{\frac{1}{\beta}(q_j - r_j)}}{e^{\frac{1}{\beta}(q_k-(r_j+r_l))} + \sum_{i\in\mathcal{O}} e^{\frac{1}{\beta}(q_i-r_i)}} \cdot \frac{\sum_{i\in\mathcal{O}} e^{\frac{1}{\beta}(q_i-r_i)}}{e^{\frac{1}{\beta}(q_j-r_j)}} \\
&= 1 - \frac{\sum_{i\in\mathcal{O}} e^{\frac{1}{\beta}(q_i-r_i)}}{e^{\frac{1}{\beta}(q_k-(r_j+r_l))} + \sum_{i\in\mathcal{O}} e^{\frac{1}{\beta}(q_i-r_i)}} \\
&= 1 - \frac{\sum_{i\in\mathcal{O}} e^{\frac{1}{\beta}(q_i-r_i)}}{e^{\frac{1}{\beta}(q_k-(r_j+r_l))} + \sum_{i\in\mathcal{O}\setminus\{k\}} e^{\frac{1}{\beta}(q_i-r_i)}}
\end{aligned}
\tag{7}
$$

The last equality in (7) holds because $\mathcal{O}\setminus\{k\} = \mathcal{O}$ if $k \notin \mathcal{O}$. Note that, given our assumption that $r_l < r_k - r_j$, we can neglect the minimum operator in equation (5) and thus, it is equal to the last line of (7). Reinserting the parameters β_{m_l} for β, j_l for j, k_l for k, $\mathcal{O}_l$ for $\mathcal{O}$ (all four $\forall l \in \mathcal{L}$), and q_{im_l} for $q_i \forall i \in \mathcal{J}, l \in \mathcal{L}$ into (7) yields Eq. (2). □

3.3 Upsell optimization problem

In this section, we formulate the upsell optimization problem as a non-linear expectation model. A first intuitive version is given by (8)–(11). The objective function (8) maximizes the additional expected revenue from upsells by optimizing prices r_l. It is subject to several constraints. In detail, constraints (9) relate to each resource h's capacity and ensure that upsells will not lead to the capacity being exceeded. The first summand of (9) counts the occupied capacity on resource h. The second summand represents the number of customers who move to resource h via an upsell from a lower resource $g < h$. The third summand represents customers leaving resource h because they buy an upsell to a higher resource $f > h$. Obviously, the second (third) summand is zero if h is the lowest (highest) resource. As no new customers are arriving during the upsell horizon, we can neglect $h = 1$, as customers only leave the lowest-valued resource. Furthermore, the prices' domain is restricted by (10). Constraints (11) simply define the upsell response probabilities p_l according to equation (2) to ease notation in (8) and (9). Note that p_l depends on r_l and, thus, the objective function (8) is non-linear.

$$
\max_{r_l} \sum_{l\in\mathcal{L}} r_l \cdot y_l \cdot p_l \tag{8}
$$

subject to

$$\sum_{l\in\mathcal{L}} y_l \cdot a_{j_l h} + \sum_{g<h}\sum_{l\in\mathcal{L}} y_l \cdot p_l \cdot a_{j_l g} \cdot a_{k_l h} - \sum_{f>h}\sum_{l\in\mathcal{L}} y_l \cdot p_l \cdot a_{j_l h} \cdot a_{k_l f} \leq c_h \qquad \forall h \in \mathcal{H}\backslash\{1\} \tag{9}$$

$$0 \leq r_l < r_{k_l} - r_{j_l} \quad \forall l \in \mathcal{L} \tag{10}$$

$$p_l = 1 - \frac{\sum_{i\in\mathcal{O}_l} e^{\frac{1}{\beta_{m_l}}(q_{im_l}-r_i)}}{\sum_{i\in\mathcal{O}_l\backslash\{k_l\}} e^{\frac{1}{\beta_{m_l}}(q_{im_l}-r_i)} + e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l}-(r_{j_l}+r_l))}} \qquad \forall l \in \mathcal{L} \tag{11}$$

This formulation effectively treats demand as if it was deterministic and equal to its expected value. Although intuitive at first glance, there is an important issue with the above model. Constraints (11) restrict upsell prices to an intuitive range. However, the upper bound on upsell prices implies a lower bound on the upsell response probability. This lower bound exceeds 0, which, in turn, may render constraint (9) infeasible, i.e., the sum of customers willing to accept an upsell exceeds the remaining capacity.

To overcome this issue, we introduce an additional decision variable u_l, which enables the company to decide directly on the number of upsells that are sold to class l, analogous to the partitioned allocation used in the standard DLP model for capacity control (see, e.g., Talluri and Van Ryzin 2004b, Chapter 3.3.1). Obviously, u_l is nonnegative and cannot exceed the expected demand for upsells. Rewriting the mathematical model including u_l results in the model (12)–(16).

$$\max_{(r_l,u_l)} \sum_{l\in\mathcal{L}} r_l \cdot u_l \tag{12}$$

subject to

$$\sum_{l\in\mathcal{L}} y_l \cdot a_{j_l h} + \sum_{g<h}\sum_{l\in\mathcal{L}} u_l \cdot a_{j_l g} \cdot a_{k_l h} - \sum_{f>h}\sum_{l\in\mathcal{L}} u_l \cdot a_{j_l h} \cdot a_{k_l f} \leq c_h \qquad \forall h \in \mathcal{H}\backslash\{1\} \tag{13}$$

$$0 \leq r_l < r_{k_l} - r_{j_l} \qquad \forall l \in \mathcal{L} \tag{14}$$

$$0 \leq u_l \leq y_l \cdot p_l \qquad \forall l \in \mathcal{L} \tag{15}$$

$$p_l = 1 - \frac{\sum_{i\in\mathcal{O}_l} e^{\frac{1}{\beta_{m_l}}(q_{im_l}-r_i)}}{\sum_{i\in\mathcal{O}_l\backslash\{k_l\}} e^{\frac{1}{\beta_{m_l}}(q_{im_l}-r_i)} + e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l}-(r_{j_l}+r_l))}} \qquad \forall l \in \mathcal{L} \tag{16}$$

This model is still non-linear in its objective function and some constraints. To solve it, we linearize the constraints by optimizing the probabilities instead of the upsell prices. This change of bijective variables is a standard trick in dynamic pricing and has been used, for example, by Dong et al. (2009). The NLP model reformulation with linearized constraints can be found in "Appendix 1".

3.4 Benchmark approach with decoupled choice situation

While Proposition 1 states that the upsell response probability is conditional on the customers' behavior in the booking horizon, a naive approach—which we will use as a benchmark in our numerical study—may neglect this information and instead assume that the upsell acceptance decision is an entirely new choice situation, statistically decoupled from the original choice. Technically, this means that the company assumes i.i.d. Gumbel-distributed random variables ϵ_{nj} and ϵ_{nk} for staying with the originally bought product j and the upsell to k, respectively. The resulting binary logit model yields the unconditional upsell response probability p_l^u given by (17). It can be used in the upsell optimization model instead of (2).

$$p_l^u = \frac{e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l} - (r_{j_l} + r_l))}}{e^{\frac{1}{\beta_{m_l}}(q_{j_l m_l} - r_{j_l})} + e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l} - (r_{j_l} + r_l))}} \tag{17}$$

The following proposition clarifies the relation of p_l^u and p_l.

Proposition 2 *The naive approach overestimates the upsell response probability, i.e., $p_l \leq p_l^u$ for $k_l \in \mathcal{O}_l$ as well as for $k_l \notin \mathcal{O}_l$.*

Proof First, we show that $p_l \leq p_l^u$ when the company has only two possible products (i.e., $|\mathcal{J}| = 2$): the bought product j_l and the product to which an upsell will be offered k_l. If $k_l \notin \mathcal{O}_l$, then obviously $p_l = p_l^u$. In this case, the company cannot learn anything regarding the customer's preferences for product k_l from the booking horizon.

If $k_l \in \mathcal{O}_l$ in the two-product case, then (2) and (17) turn into

$$p_l = 1 - \frac{e^{\frac{1}{\beta_{m_l}}(q_{j_l m_l} - r_{j_l})} + e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l} - r_{k_l})}}{e^{\frac{1}{\beta_{m_l}}(q_{j_l m_l} - r_{j_l})} + e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l} - (r_{j_l} + r_l))}} \quad \text{and} \quad p_l^u = \frac{e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l} - (r_{j_l} + r_l))}}{e^{\frac{1}{\beta_{m_l}}(q_{j_l m_l} - r_{j_l})} + e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l} - (r_{j_l} + r_l))}}.$$

We compare these two probabilities as follows:

$$
\begin{aligned}
p_l &= 1 - \frac{e^{\frac{1}{\beta_{m_l}}(q_{j_l m_l}-r_{j_l})} + e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l}-r_{k_l})}}{e^{\frac{1}{\beta_{m_l}}(q_{j_l m_l}-r_{j_l})} + e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l}-(r_{j_l}+r_l))}} \\
&= \frac{e^{\frac{1}{\beta_{m_l}}(q_{j_l m_l}-r_{j_l})} + e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l}-(r_{j_l}+r_l))} - e^{\frac{1}{\beta_{m_l}}(q_{j_l m_l}-r_{j_l})} - e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l}-r_{k_l})}}{e^{\frac{1}{\beta_{m_l}}(q_{j_l m_l}-r_{j_l})} + e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l}-(r_{j_l}+r_l))}} \\
&= \frac{e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l}-(r_{j_l}+r_l))} - e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l}-r_{k_l})}}{e^{\frac{1}{\beta_{m_l}}(q_{j_l m_l}-r_{j_l})} + e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l}-(r_{j_l}+r_l))}} \\
&= \frac{e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l}-(r_{j_l}+r_l))}}{e^{\frac{1}{\beta_{m_l}}(q_{j_l m_l}-r_{j_l})} + e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l}-(r_{j_l}+r_l))}} - \frac{e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l}-r_{k_l})}}{e^{\frac{1}{\beta_{m_l}}(q_{j_l m_l}-r_{j_l})} + e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l}-(r_{j_l}+r_l))}} \\
&\le \frac{e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l}-(r_{j_l}+r_l))}}{e^{\frac{1}{\beta_{m_l}}(q_{j_l m_l}-r_{j_l})} + e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l}-(r_{j_l}+r_l))}} = p_l^u.
\end{aligned}
$$

The last inequality holds because the numerators as well as the denominators are always positive ($e^x > 0 \,\forall\, x \in \mathbb{R}$). Up to this point, we showed $p_l \le p_l^u$ for the two-product case.

To generalize this statement, it is sufficient to show that by adding more products, p_l cannot increase—and thus, to show that the fraction in formula (1) cannot decrease. Adding other products than j_l and k_l to the offer set corresponds to adding a positive constant $\delta = \sum_{i \in \mathcal{O}_l \setminus \{j_l, k_l\}} e^{\frac{1}{\beta_{m_l}}(q_{i m_l}-r_i)}$ to both the numerator and the denominator. This is always bigger than the original term:

$$
\frac{e^{\frac{1}{\beta_{m_l}}(q_{j_l m_l}-r_{j_l})} + e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l}-r_{k_l})} + \delta}{e^{\frac{1}{\beta_{m_l}}(q_{j_l m_l}-r_{j_l})} + e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l}-(r_{j_l}+r_l))} + \delta} \ge \frac{e^{\frac{1}{\beta_{m_l}}(q_{j_l m_l}-r_{j_l})} + e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l}-r_{k_l})}}{e^{\frac{1}{\beta_{m_l}}(q_{j_l m_l}-r_{j_l})} + e^{\frac{1}{\beta_{m_l}}(q_{k_l m_l}-(r_{j_l}+r_l))}}.
$$

Consequently, p_l decreases with any increasing δ. As p_l^u does not depend on the availability of other products in the offer set, and thus, remains constant, $p_l \le p_l^u$ is still valid for $k_l \in \mathcal{O}_l$ and $k_l \notin \mathcal{O}_l$. Thus, overall, p_l^u overestimates the true upsell response probability p_l. □

As a consequence of Proposition 2, for the optimization model (12)–(16), using the overestimating probability p_l^u instead of p_l implies higher prices to comply with the constraints. As upsell prices are optimal with p_l, the expected revenues by upsells are likely to decrease with p_l^u.

3.5 States of information

When presenting the optimization model above, we assumed that the company recorded all the necessary data during the selling period, i.e., the company knows for each customer

- the bought product,
- the offer set from which the customer chose, and
- her segment m, i.e., the deterministic part of her utility defined by the quality indices q_{jm} and the scale parameter β_m.

According to our experience from practice, this assumption is not always fulfilled. While knowing the bought product is a prerequisite for making meaningful targeted upsell offers—and is natural to record—the two other pieces of information may be missing. Thus, we consider the following three alternative states of information that can occur (see also Fig. 3):

- SOI 1: Full information on bought product, offer set, and customer segment,
- SOI 2: Only bought product and customer segment are known,
- SOI 3: Only the bought product is known.

In the case of SOI 2 and SOI 3, the missing information needs to be imputed. Here, instance-specific approaches seem natural, and we will investigate a straightforward one in our numerical experiments (see Sect. 4).

3.6 Group bookings

In the case of group bookings, i.e., customers traveling together also buy their tickets together, it is very likely that they want to sit together in one compartment. Thus, these customers accept the upselling offer if and only if all group members are willing to accept it. Let θ_v be the fraction of customers belonging to a group of v persons and $\theta_1 = 1 - \sum_{v>1} \theta_v$, then the upsell response probability for a customer class l is

$$p_l^{\mathrm{gr}} = \sum_v \theta_v \cdot p_l^v \qquad \forall l \in \mathcal{L} \tag{18}$$

with p_l being calculated by formula (2) for each group member. Formula (18) can then be used in the upsell optimization model instead of (2).

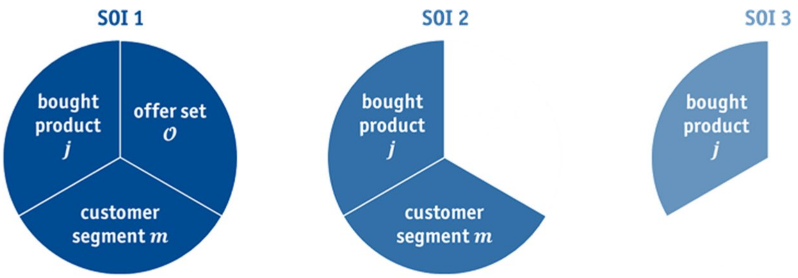

Fig. 3 States of information: Level of detail in the recorded booking data

4 Numerical study

In this section, we numerically analyze the benefit of upsell price optimization and using data from the booking horizon (selling period) with an example from the airline industry. First, we introduce the test instances (Sect. 4.1). In Sect. 4.2, we present the imputation procedure if information from the booking horizon is missing (SOI 2 and SOI 3). Then, we describe the upsell price optimization approaches and benchmarks investigated (Sect. 4.3). Section 4.4 details the evaluation procedure. In Sect. 4.5, we present the numerical results.

4.1 Test instances

We consider a single flight with three compartments (resources) ($\mathcal{H} = \{1 = economy, 2 = business, 3 = first\}$) and capacities of $c = (60, 30, 10)$ seats. In each compartment, two fares are defined (see Fig. 4) that may differ in the associated restrictions so that it makes sense to simultaneously offer them. For illustration purposes, think of a cheap saver fare (e.g., without rebooking and cancellation) and a more expensive full fare that offers flexibility. The saver fare products 1, 3, and 5 are priced at €50, €230, and €450, respectively. The full-fare products 2, 4, and 6 are priced at €110, €285, and €560, respectively.

Two customer segments are considered. Leisure customers (segment $m = 1$) do not need the full fare products' additional flexibility and, thus, have the same quality indices for both fares in a compartment. In detail, we have $q_{11} = q_{21} = 75$ for economy class, $q_{31} = q_{41} = 225$ for business class, and $q_{51} = q_{61} = 430$ for first class. Business customers (segment $m = 2$) require the flexibility and, thus, have quality indices of −70, 120, 0, 300, 0, and 555. The logit model's scale parameters are $\beta_1 = 20$ and $\beta_2 = 15$. The no purchase utilities q_0m reflect overall demand intensity and differ across four instances that we are going to consider: Instance I1 models high demand with $q_0 = (-25, -15)$, instance I2 intermediate demand with $q_0 = (0, 0)$, instance I3 low demand with $q_0 = (12.5, 7.5)$, and instance I4 very low demand with $q_0 = (25, 15)$. The rationale for the choice of parameters is given in "Appendix 2".

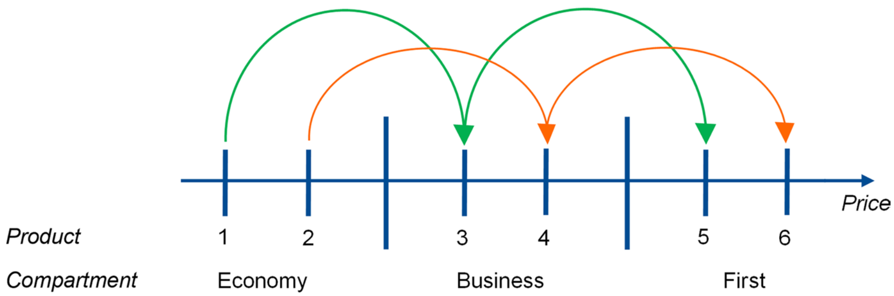

Fig. 4 Products and possible upsells

The products are sold during a booking horizon of $T = 150$ time periods. In each time period, with probability $\lambda_1 = 0.5$ ($\lambda_2 = 0.3$), a leisure (business) customer arrives. During the booking horizon, the airline (company) uses a basic capacity control approach. In detail, at the beginning of the booking horizon, a choice-based linear program (CDLP, see Gallego et al. 2004) is solved. Whenever a customer arrives, the offer set is randomly selected according to the CDLP's primal solution. To emulate unforeseen demand/capacity mismatches, the booking process abruptly ends after 80% of the horizon in period $t = 120$. After the booking horizon, the upsell horizon begins.

As common in practice, we only offer upsells to the next higher compartment—analogous to limited cascading upgrading (see Steinhardt and Gönsch 2012). Airlines often rather leave first-class seats empty than upgrade an economy passenger to it. Additionally, customers stay with the same fare type (saver or full fare). Thus, for economy and business class products, exactly one upsell possibility exists. The resulting upsell hierarchy is illustrated in Fig. 4. The parameters are summarized in Table 5 in the "Appendix 3".

4.2 Imputation procedure

In SOI 1, the airline knows for every customer the product she bought, her segment, and her offer set (see also Sect. 3.5 for a detailed description of the three SOI). Thus, no imputation is necessary. In SOI 2, the airline only knows the bought product and the customer segment, but missed recording the offer set, which often happens in practice. We handle this deficit as follows in our approaches: if the customer bought a saver fare product, we assume that the offer set consisted of the three saver products. Analogously, if a full-fare ticket was bought, we assume all three full-fare products were offered. In SOI 3, only the bought product is known. To handle this deficit, analogously to SOI 2 above, we deduce the offer set from the bought product. Likewise, we assume that a customer who bought a full (saver) fare product belongs to the business (leisure) segment.

4.3 Upsell price optimization approaches

The following upsell price optimization approaches and benchmarks are investigated:

- *C-CP* is the Customer class-specific price setting with Conditional Probabilities. This is the new approach using the optimization model (12)–(16) to determine prices tailored to each customer class. It exploits the fact that, from customers' choices during the booking process, we obtain information and now have a conditional distribution regarding their stochastic utility.
- *S-CP* is an intuitive benchmark that sets prices according to a single static Share with Conditional Probabilities. Instead of optimizing one upsell price for each customer class, we restrict the optimization to one single decision variable: the discount of an upsell compared to the difference of the associated products'

prices. More precisely, the airline decides on a global share s of the difference between the price of the bought product j_l and the upsell product k_l. Mathematically, this is $r_l = s \cdot (r_{k_l} - r_{j_l})\ \forall\, l \in \mathcal{L}$. Technically, we still use model (12)–(16) and substitute r_l by the given formula.

- *C-UP* is the benchmark approach with the decoupled choice situation as described in Sect. 3.4, i.e., with Customer class-specific price setting assuming the Unconditional Probability. More precisely, the airline neglects what it learned about customers' stochastic utility during the booking horizon. Technically, it does not condition customers' stochastic utility on their choice during the booking horizon but still assumes Gumbel distributed error terms $\in_{nj}$. Thus, instead of the conditional probability (16), the unconditional binary logit model (17) is used.

4.4 Evaluation procedure

To evaluate the pricing methods, in each setting, we first simulate $N = 10.000$ realizations of the booking horizon for each of the four instances and use them for all SOI and upsell price optimization approaches, in the sense of common random numbers.

In each of these simulation runs, the airline performs capacity control using the CDLP's primal solution as described above. Customers sequentially arrive and choose according to a MNL from the offer set presented to them. Thus, at the end of the booking horizon, the flight has a certain capacity of remaining free seats in the three compartments. The customers who bought a ticket are partitioned into customer classes as described in Sect. 3.1 (combinations of customer segment, seen offer set, and bought product).

For the subsequent upsell decision-making, the size of each resulting customer class as well as the remaining overall capacity are relevant. For SOI 1, customer information is directly used within the upsell price optimization approaches to determine upsell prices. For SOI 2 and SOI 3, customer information is degraded as described above, the missing information is imputed, and fed into the optimization approaches.

Finally, the resulting upsell prices are evaluated. To do so, we use model (12)–(16) with full information (SOI 1) as this is supposed to reflect the actual situation in reality. Consequently, we also use this model to evaluate prices obtained with approaches for SOI 2 and SOI 3, and also when upgrade decisions have been made on the (simplifying) assumption of unconditional probabilities by our benchmark C-UP.

4.5 Results

The following results are averages over all simulation runs, ignoring runs in which business and first class are both sold out at the end of the booking horizon (about

7–17% depending on the instance) such that no upselling would be possible. Table 1 gives an overview of the situation at the end of the booking horizon.

4.5.1 C-CP

The first method analyzed is C-CP. Figure 5 shows the expected upsell revenue from the evaluation (second bar in each group of bars) for each instance relative to the upsell revenue obtained in the SOI 3 evaluation. Broadly speaking, we see that the revenue usually decreases in the state of information. More precisely, in SOI 2 and SOI 3 we obtain always significantly less expected revenue than in SOI 1.

Compared to SOI 3, where the least data is recorded during the booking horizon, SOI 1 allows 2–5% higher revenues. Note that the very slight advantage of the evaluation of SOI 3 over SOI 2 in instance I1 is instance-specific and may result from the fact that many leisure customers are misclassified as business customers in SOI3 and are thus assigned higher quality indices. This effect in turn can lead to lower upsell prices. In addition, the customer classes can get larger which can result in more customers receiving an upsell offer.

Additionally, in Fig. 5, the first bars in each group of bars show the objective function value of the optimization model used to obtain the prices for each instance and SOI. For SOI 1, this value is identical to the expected revenue by definition. For SOI 2 and SOI 3, the objective function values are based on the assumption that the imputation is perfect. As it is not, they tend to overestimate the revenue.

Regarding the general benefit from upselling, it can be stated that SOI 1 yields the highest additional revenue in all instances compared to the revenue from the respective booking horizon. Revenues increase up to 15.80% for instance I1 and up to 16.01%, 16.81%, and 13.97% for I2, I3, and I4, respectively (not shown in the figure).

To demonstrate significant differences in the methods, Table 2 shows the confidence intervals (in absolute values) based on a paired t-test with a significance level of $\alpha = 0.99$ (see the corresponding confidence intervals for all instances and customer class-specific pricing methods in Table 2 and Tables 6, 7 and 8 (Appendix 3)). Here again, we can see that SOI 1 always outperforms SOI 2 and SOI 3. This even holds for instance I4, where the revenues in SOI 1–3 are close to each other. The evaluation values of the methods in the column of each of these tables are subtracted from the ones of the methods in the corresponding row. Consequently, positive (negative) mean values in the tables demonstrate a dominance of the methods in the corresponding row (column).

To analyze the potential benefits of upselling in a more systematic way, we now partition the set of simulation runs into subclasses according to the free capacity observed after the booking horizon. We define eight subclasses by distinguishing free capacity in the three compartments below/above its median value. For instance I2, Table 3 shows the subclasses together with their expected upsell revenue relative to the revenue from the booking horizon. Results for the other instances are similar and included in Table 9 in the “Appendix 3”.

We observe that more free capacity in business and first class leads to more revenue and that free seats in the first-class compartment generate more revenue

Table 1 Situation at the end of the booking horizon (average values)

	I1—high demand	I2—intermediate demand	I3—low demand	I4—very low demand
Revenue in booking horizon	14,044.87	13,322.09	12,400.06	10,812.23
Number of runs where upsells are possible	8330	8866	8763	9384
Remaining capacity (*economy*, *business*, *first*)	(11.57, 6.5, 2.38)	(12.01, 6.26, 2.45)	(25.48, 6.39, 2.45)	(49.49, 6.28, 2.45)

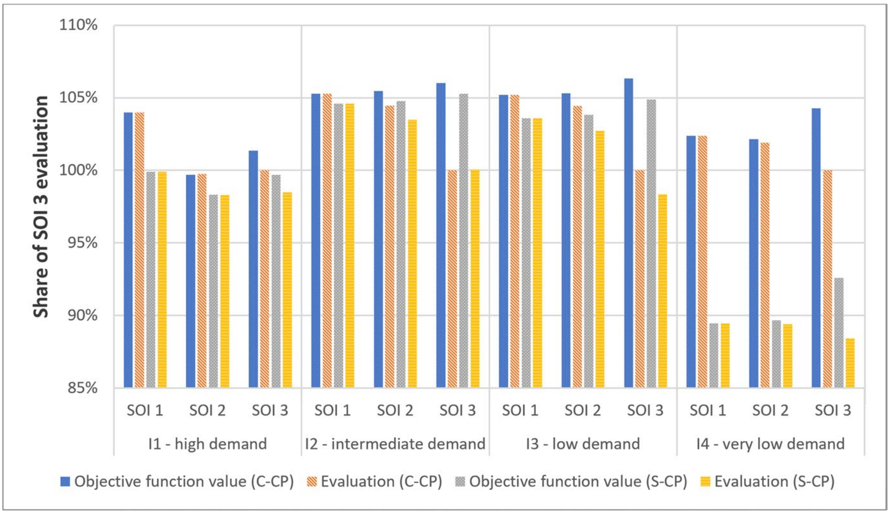

Fig. 5 Expected upsell revenue (C-CP and S-CP) relative to revenue in SOI 3. Objective function values refer to the respective SOI, evaluation values are retrieved by using prices from the respective SOI, but evaluating on SOI 1

Table 2 Confidence intervals of the differences of the evaluation values for Instance I1

	C-CP, SOI 1	C-CP, SOI 2	C-CP, SOI 3	C-UP, SOI 1	C-UP, SOI 2	C-UP, SOI 3
C-CP, SOI 1	–					
C-CP, SOI 2	85.97 ± 0.78	–				
C-CP, SOI 3	81.08 ± 1.07	− 4.89 ± 1.02	–			
C-UP, SOI 1	1118.65 ± 13.61	1032.68 ± 13.22	1037.57 ± 13.29	–		
C-UP, SOI 2	1230.58 ± 12.79	1144.60 ± 12.43	1149.50 ± 12.48	111.93 ± 7.05	–	
C-UP, SOI 3	1470.95 ± 14.95	1384.98 ± 14.52	1389.87 ± 14.57	352.30 ± 6.41	240.37 ± 5.8	–

Positive mean values are interpreted such as that the method in the column is better than the one in the corresponding row

Table 3 Remaining capacity after booking horizon and SOI for each subclass

Subclass	Remaining capacity after booking horizon			SOI		
	Eco	Business	First	1	2	3
1	< 12	< 6	< 2	7.30	7.24	6.92
2	< 12	< 6	> 2	19.15	19.06	18.33
3	< 12	> 6	< 2	14.80	14.66	13.87
4	< 12	> 6	> 2	28.95	28.71	27.55
5	> 12	< 6	< 2	6.39	6.34	6.07
6	> 12	< 6	> 2	17.31	17.22	16.60
7	> 12	> 6	< 2	13.63	13.50	12.81
8	> 12	> 6	> 2	26.60	26.37	25.37

than free business seats. This is intuitive as a free first-class seat always ‘includes’ a free business seat (given customers strictly prefer higher compartments). It offers the opportunity for two upgrades: First, from the business class compartment to first class, and then a second one from economy that uses the business class seat which just freed up.

Of course, it is also important how many passengers are eligible for an upsell offer. This is the reason why the situation is different for the economy class. Comparing subclasses 1–5, 2–6, 3–7, and 4–8, we see that the more sales (and, thus, the fewer free seats) we have in economy class, the higher the expected revenue is. This effect also seems to cause the decrease in revenue with decreasing demand, that is, from I1 to I4 (see Table 9 in the “Appendix 3” in combination with average revenues from the booking horizon of the instances in Table 1). This highlights again the upsell revenue’s dependency on economy ticket sales because the average number of free seats in business and first class is nearly the same over all four instances.

4.5.2 S-CP

The first benchmark approach analyzed is the static share S-CP. As expected, constraining the solution space leads to a loss in revenue (fourth bars in each group in Fig. 5). The losses can be severe and range from 2 to 14% (compare Evaluation C-CP and S-CP in Fig. 5). The differences of the objective function values and evaluation values from S-CP show similar behavior as from C-CP. Apparently, these differences are mainly driven by the SOI.

4.5.3 C-UP

Finally, we analyze what happens if the airline does not learn from the booking horizon, that is, it does not condition customers’ stochastic utilities on their choices. Figure 6 shows the objective function value and expected revenue obtained with C-UP relative to the corresponding values obtained with C-CP. The objective function values (left bars in each group of bars) are between 109 and 114%. A value greater than 100% implies that the upsell response probability was overestimated, leading to a higher objective value, thus confirming our analytical result from Proposition 2. More precisely, the unconditional probability according to the binary logit model suggests that more customers are willing to buy an upsell at the price selected than actually are. This, in turn, leads to upsell prices set higher than in the C-CP, which causes huge revenue losses during the upsell horizon, as shown in the second bar in each group. For example, with the best information on demand (SOI1), C-UP’s upsell revenue only reaches about 13–45% of C-CP’s revenue.

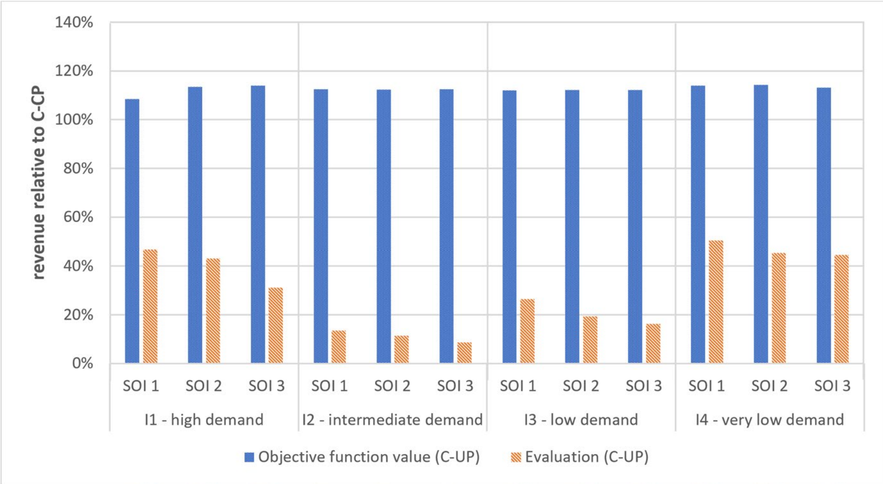

Fig. 6 Objective function value and expected revenue from C-UP relative to C-CP

4.5.4 Group bookings

Up to now, each customer decided individually whether she accepts an upsell offer. In reality, this might not be true: Imagine a couple travelling together. Then, they will only accept the upsell if both accept the upsell offer. The probability that two persons accept an upsell is generally lower than the one for one customer alone (see 3.6). First, to investigate the revenue loss resulting from not considering group bookings in the upsell price optimization problem, we optimize on single customers and evaluate on a 50% ratio of group bookings. This means that the airline assumes that all customers travel alone, but actually 50% of the passengers travel in groups. Technically, the prices from the optimization approach C-CP (see Sects. 4.3 and 4.5.1) are evaluated with groups (see the second bar of each scenario in Fig. 7). For comparison, Fig. 7 additionally shows the mean C-CP evaluation values as first bar of each scenario. Next, we correctly consider the ratio of groups in the optimization problem (i.e., the airline assumes the 50% ratio already in their optimization). The result can also be found in Fig. 7 (third bars). As expected, using correct data for groups in the optimization generally outperforms assuming single customers in the presence of groups. In Instance I1, the improvement of the second over the third bar in SOI 2 does not seem intuitive, but can be explained by smaller assumed offer sets in the optimization problem. The smaller the offer set, the more decreases the upsell response probability. If it decreases, prices are set lower, and, consequently, the probability of groups accepting an upsell increases (see Eq. (18)).

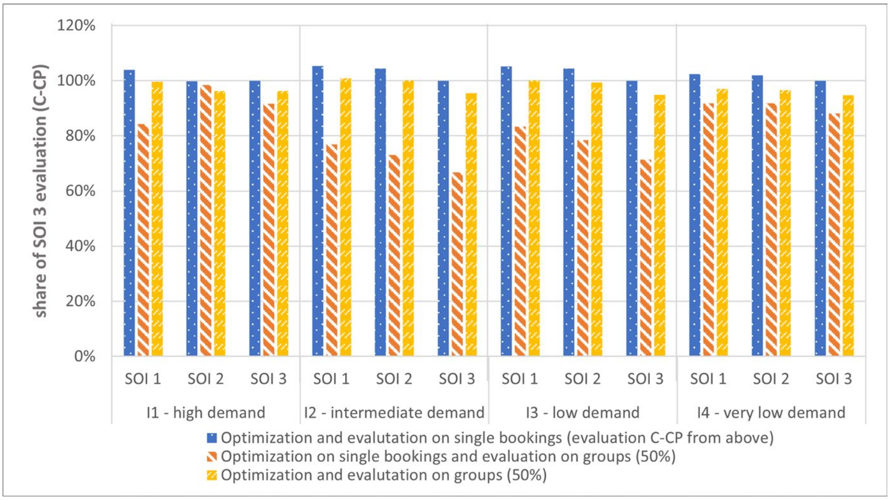

Fig. 7 Expected upsell revenue with 50% group bookings relative to revenue in SOI 3 with single bookings

4.5.5 Naive customer segment forecast

In the previous subsections, we assumed that the customer segment is either known (SOI 1 and SOI 2), or is determined by considering the bought product (SOI 3, for more details, see Sect. 4.2). The idea of the naïve customer segment forecast is that the airline does not know the original customer segments, but the airline knows the customer segments' distribution. We then approach the situation as follows. The arrival probabilities are already known in the booking horizon, e.g., from a market study. For the upsell horizon, these arrival probabilities are scaled to sum up to one. We then randomly assign customer segments to the customers (according to the normed arrival probabilities) by means of the inverse transformation method. For the upsell optimization problem, we classify the customers with their new characteristics into classes $l \in \mathcal{L}$. Compared to the original data, the naïve forecasting method misclassifies 37.5% (62.5%) of leisure (business) customers as business (leisure) customers in each of the instances. The resulting upsell prices are, obviously, evaluated again using the original customer data (SOI 1).

Table 4 Confidence intervals for the differences of the evaluation values of C-CP (SOI 1) and of the method with a naive customer segment forecasting

Instance I1	Instance I2	Instance I3	Instance I4
187.79 ± 3.22	201.77 ± 4.00	221.79 ± 4.48	290.62 ± 4.30

The results of the comparison to the SOI 1 (full information) via a paired t-test on the evaluation values of C-CP (SOI 1) and the above-mentioned method in Table 4 show, that the naive forecast method performs not only worse than SOI 1 but also worse than SOI 2 or SOI 3 (both C-CP, see Tables 2, 6, 7 and 8). Note that the latter observation is always dependent on the selected imputation procedure and on the instance. Lower error probabilities or worse imputation procedures for SOI 2 and SOI 3 could cause the naive customer segment forecast to outperform SOI 2 and SOI 3.

5 Managerial insights and conclusion

In this paper, we propose an expectation model to optimize upselling. We model customer behavior according to an MNL model and derive a closed-form expression for the upsell response probability based on customer-specific booking data from the booking horizon and the MNL. Specifically, we assume that customers' preferences, as they are modeled by their stochastic utility, stay constant over time.

A numerical study generated numerous managerial insights. First, recording detailed disaggregated booking data is indeed recommended, because this data can be used to considerably increase revenue from upsells. We observed mean increases of up to 5.28% in the study. If not all the required data is available, such as the original offer set seen by the customer or segment information, it must be imputed. This 'guessing' is shown to often have a very negative impact on revenues. Second, targeting upsells to finer customer classes pays off. While offering all upsells at a common discount already increases revenue, the benefit is much larger with upsell prices tailored to customer classes' preferences. Third, the upsell revenue obviously increases with the number of customers that can be upgraded. For example, the potential is higher with many economy bookings and few in business class. Fourth, we saw that knowledge about how customer preferences evolve over time (i.e., between the time of purchase and the upsell offer) is important. We assumed them to be stable and observed that using a benchmark approach with preferences that are uncorrelated over time—implying that the customer's upsell choice situation is uncoupled from the original one—led to huge revenue losses. Fifth, knowledge about group bookings is essential in the optimization process, because customers in groups behave differently to single booking customers. Groups are less likely to accept an upsell, and not adequately considering them may lead to high revenue losses.

There are several potential avenues for further research. First, regarding the evolution of preferences over time mentioned above, we focused on two extremes: no change and no correlation over time. It seems desirable to also capture the continuum in between. Second, customer choice may follow another model than the MNL model, e.g., nested logit models (see e.g. Train 2009, Chapter 4.2), the general attraction model (Gallego et al. 2015), or random distributions with given probability function. If closed-form solutions exist, they could be integrated into the upsell optimization model. Otherwise, different approaches seem necessary. Third, more sophisticated imputation methods for the different states of information tailored to specific application areas could be investigated. Fourth, similar to dynamic pricing with strategic customers, the firm and/or the customers could be forward-looking in the sense that they maximize utility over time and already consider during the booking horizon the possibility of later upsell offers. Fifth, upselling can be combined with different practices like overbooking or frequent flyer programs to investigate the combined effects. In the case of overbooking of the lower compartments, but not the overall capacity, our upsell optimization model can already be used as it is to upsell customers.

Appendix 1: Model with linearized constraints

Replace r_l in (12)–(16) by solving (16) for r_l to get the model (19)-(22).

$$\max_{(p_l,u_l)} \sum_{l\in\mathcal{L}} u_l \cdot \left(\beta_l \cdot \ln(1-p_l) - \beta_l \cdot \ln\left(\left(\left(e^{\frac{q_{0l}}{\beta_l}} + \sum_{i\in\mathcal{O}_l} e^{\frac{1}{\beta_l}(q_{li}-r_i)}\right) - (1-p_l) \cdot \left(e^{\frac{q_{0l}}{\beta_l}} + \sum_{i\in\mathcal{O}_l\setminus\{k_l\}} e^{\frac{1}{\beta_l}(q_{li}-r_i)}\right)\right)\right) + q_{lk_l} - r_{j_l}\right) \tag{19}$$

subject to

$$\sum_{l\in\mathcal{L}} y_l \cdot a_{j_l h} + \sum_{g<h} \sum_{l\in\mathcal{L}} u_l \cdot a_{j_l g} \cdot a_{k_l h} - \sum_{f>h} \sum_{l\in\mathcal{L}} u_l \cdot a_{j_l h} \cdot a_{k_l f} \le c_h \qquad \forall h \in \mathcal{H} \tag{20}$$

$$0 < p_l < 1 \qquad \forall l \in \mathcal{L} \tag{21}$$

$$0 \le u_l \le y_l \cdot p_l \qquad \forall l \in \mathcal{L} \tag{22}$$

Appendix 2: Rationale for choice parameters

We want to provide the rationale for the choice of parameters in the numerical study to facilitate further research on upselling. Therefore, we first present required values and then papers with similar values. Last, we show the connection of our values to the ones given in Gallego et al. (2009).

Required values

In terms of capacity, we aimed for a realistic aircraft size. Many papers use smaller capacities [see, e.g., Gallego et al. (2009) with up to 24 seats on a single leg, Bront et al. (2009) with between 30 and 50 seats per leg, 5–10 seats/leg in Liu and Van Ryzin (2008)]. For modelling customer choice behavior, we need the quality index q, the price for each product r_j, and the scale parameter β.

With these, we can calculate the deterministic utility $V_j = q_j - r_j$ or the preference weight $v = e^{(\frac{V}{\beta})}$. Most papers directly define these parameters instead of breaking them down into the parameters that we need for our approach. An exemplary paper with deterministic utilities is Wang et al. (2016). For preference weights, please refer to Liu and Van Ryzin (2008) and Bront et al. (2009).

Papers with similar values

The most similar type of parameters for a numerical study can be found in Gallego et al. (2009). The authors define the deterministic utility as the difference between schedule quality and price (both weighted). Although this can be adapted to our model, i.e., using fare quality (quality indices for flexibility and saver fare) instead of schedule quality, we still see problems in the adaption. First, the authors do not consider different customer segments. This forces us to come up with different values for the second customer segment, which, in turn, makes it very difficult to compare the instances. Second, the authors consider different flights and not different compartments. Thus, we would have to come up with quality indices for the different compartments. Another paper with a similar study is Dong et al. (2009). Here again, we do not have the full parameter set for the example in our manuscript. Although the authors provide quality indices (a) and cost (c), they do not numerically define their scaling parameter μ for the MNL model. Thus, with a variable scaling parameter, we are able to match one of the given instances with any product pair. This already implies the second adaption problem: we would have to propose choice model values for four additional products as they only consider two products in their numerical study.

Altogether, therefore, there appears to be no paper whose values we can use directly. Of course, we could use some values from literature (like capacity and choice parameters for one product), but we would still have to come up with a multitude of other parameters. This, frankly, would create a largely new setting anyway, which cannot be attributed to the literature. Nonetheless, we oriented our values on values from literature. More precisely, we have chosen r, q, and β such that the absolute value ranges of $\frac{1}{\beta} \cdot (q - r)$ are similar to the ones in Gallego et al. (2009).

Appendix 3: Tables

See Tables 5, 6, 7, 8, 9

Table 5 Overview of notation and values for the numerical study

Variable	Name	Value in numerical study
$\mathcal{J}$, Indices j,k	Set of products	$\{1, 2, 3, 4, 5, 6\}$
$\mathcal{H}$, Index h, g, f	Set of resources	$\{1, 2, 3\}$ resp. $\{economy, business, first\}$
$\mathcal{M}$, Index m	Set of customer segments in booking horizon	$\{1, 2\}$
$\mathcal{N}$, Index n	Set of customers in one customer stream	
$\mathcal{L}$, Index $l = 1 \ldots, L$	Set of customer classes in upsell horizon	Depends on instance and state of information
q_{0m}	No-purchase utility	Differs among instances to scale demand
q_{jm}	Quality indices	$\begin{pmatrix} 75 & -70 \\ 75 & 120 \\ 225 & 0 \\ 225 & 300 \\ 430 & 0 \\ 430 & 555 \end{pmatrix}$
λ_m	Arrival probability	(0.5, 0.3)
r	Revenues	(0, 50, 110, 230, 285, 450, 560)
c	Capacity	(60, 30, 10)
T	Number of time periods	150
β_m	Scale parameter	(20, 15)
y_l	Number of customers in segment l	Depends on customer stream and state of information
N	Number of iterations in simulation	10,000
u_l	Decision variable, number of upsells planned	
w_{jn}	Willingness to pay	$w_{jn} = q_j + \beta\epsilon_{jn}$
V_{jm}	Deterministic utility of a customer segment m for product j	$V_{jm} = \frac{1}{\beta_m}(q_{jm} - r_j)$
ϵ_{jn}	Deterministic utility of a customer n for product j	Random Gumbel-distributed values

Table 5 (continued)

Variable	Name	Value in numerical study
$A_{\lvert\mathcal{J}\rvert\times\lvert\mathcal{H}\rvert}$	Capacity consumption	$\begin{pmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{pmatrix}$

Table 6 Confidence intervals of the differences of the evaluation values for Instance I2

	C-CP, SOI 1	C-CP, SOI 2	C-CP, SOI 3	C-UP, SOI 1	C-UP, SOI 2	C-UP, SOI 3
C-CP, SOI 1	–					
C-CP, SOI 2	15.46 ± 0.46	–				
C-CP, SOI 3	100.82 ± 1.95	85.36 ± 1.82	–			
C-UP, SOI 1	1738.12 ± 18.77	1723.46 ± 18.63	1638.10 ± 17.85	–		
C-UP, SOI 2	1783.12 ± 18.51	1767.66 ± 18.38	1682.30 ± 17.62	44.20 ± 4.65	–	
C-UP, SOI 3	1845.71 ± 18.40	1830.25 ± 18.26	1744.89 ± 17.42	106.80 ± 5.45	62.60 ± 4.73	–

Positive mean values are interpreted such as that the method in the column is better than the one in the corresponding row

Table 7 Confidence intervals of the differences of the evaluation values for Instance I3

	C-CP, SOI 1	C-CP, SOI 2	C-CP, SOI 3	C-UP, SOI 1	C-UP, SOI 2	C-UP, SOI 3
C-CP, SOI 1	–					
C-CP, SOI 2	13.76 ± 0.85	–				
C-CP, SOI 3	96.20 ± 1.94	82.44 ± 1.80	–			
C-UP, SOI 1	1426.58 ± 15.96	1412.82 ± 15.83	1330.38 ± 15.34	–		
C-UP, SOI 2	1570.38 ± 15.16	1556.62 ± 15.15	1474.18 ± 14.67	143.80 ± 8.52	–	
C-UP, SOI 3	1643.19 ± 15.62	1629.42 ± 15.58	1546.98 ± 14.97	216.60 ± 8.30	72.80 ± 5.08	–

Positive mean values are interpreted such as that the method in the column is better than the one in the corresponding row

Table 8 Confidence intervals of the differences of the evaluation values for Instance I4

	C-CP, SOI 1	C-CP, SOI 2	C-CP, SOI 3	C-UP, SOI 1	C-UP, SOI 2	C-UP, SOI 3
C-CP, SOI 1	–					
C-CP, SOI 2	6.54 ± 0.15	–				
C-CP, SOI 3	33.21 ± 0.94	26.67 ± 0.89	–			
C-UP, SOI 1	704.38 ± 10.81	697.84 ± 10.69	671.17 ± 10.46	–		
C-UP, SOI 2	778.72 ± 10.66	772.18 ± 10.54	745.51 ± 10.50	74.34 ± 3.80	–	
C-UP, SOI 3	802.21 ± 10.85	795.67 ± 10.74	769.00 ± 10.67	97.83 ± 3.81	23.49 ± 1.03	–

Positive mean values are interpreted such as that the method in the column is better than the one in the corresponding row

Table 9 Additional revenue by upsells [%] of subclasses for instances I1, I3, and I4

Instance	I1	I3	I4	I1, I3, I4		I1 - high			I3 - low			I4 - very low		
	Remaining capacity after booking horizon					SOI			SOI			SOI		
Subclass	Economy			Busi- ness	First	1	2	3	1	2	3	1	2	3
1	< 12	< 26	< 50	< 6	< 2	7.30	6.85	6.85	7.31	7.28	6.86	7.02	6.99	6.69
2	< 12	< 26	< 50	< 6	> 2	18.68	17.94	18.00	19.88	19.77	19.01	19.44	19.32	19.01
3	< 12	< 26	< 50	> 6	< 2	14.55	13.94	14.01	15.23	15.11	14.25	11.58	11.55	11.26
4	< 12	< 26	< 50	> 6	> 2	27.73	26.79	26.87	30.03	29.83	28.80	24.85	24.72	24.44
5	> 12	> 26	> 50	< 6	< 2	6.61	6.22	6.22	6.76	6.73	6.37	5.90	5.88	5.71
6	> 12	> 26	> 50	< 6	> 2	17.36	16.70	16.72	18.45	18.34	17.72	16.46	16.35	16.19
7	> 12	> 26	> 50	> 6	< 2	13.50	12.94	13.01	14.12	13.97	13.25	8.12	8.11	7.97
8	> 12	> 26	> 50	> 6	> 2	26.54	25.70	25.62	28.38	28.07	27.29	19.98	19.86	19.69

Acknowledgements We thank the anonymous reviewers for their valuable comments and suggestions, which helped us to improve the quality of the manuscript. Jochen Gönsch thanks the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) for their support of this work by the grant "Design and Control of Incompletely Specified Products" (GO 2254/1-1).

Funding Open Access funding enabled and organized by Projekt DEAL.

Data availability The data and code that supports the findings of this study are available at https://doi.org/10.7910/DVN/7162GJ.

Declarations

Conflict of interest The authors have no competing interests to declare that are relevant to the content of this article.

References

Ahn JJ, Choi EKC, Joung HW (2022) Promoting hotel upselling: the effect of message appeal and delivery setting on consumer attitude and purchase intention. J Hosp Tour Manag 52:295–303. https://doi.org/10.1016/j.jhtm.2022.07.009

Aydin G, Ziya S (2008) Pricing promotional products under upselling. Manuf Serv Oper Manag 10(3):360–376. https://doi.org/10.1287/msom.1070.0187

Belobaba PP (1987) Survey paper—airline yield management an overview of seat inventory control. Transp Sci 21(2):63–73. https://doi.org/10.1287/trsc.21.2.63

Ben-Akiva ME, Lerman SR (1985) Discrete choice analysis: theory and application to travel demand, vol 9. MIT Press, Cambridge

Bergemann D, Bonatti A, Haupt A et al (2022) The optimality of upgrade pricing. Springer, Berlin, pp 41–58. https://doi.org/10.1007/978-3-030-94676-0_3

Bitran G, Caldentey R (2003) An overview of pricing models for revenue management. Manuf Serv Oper Manag 5(3):203–229. https://doi.org/10.1287/msom.5.3.203.16031

Biyalogorsky E, Gerstner E, Weiss D et al (2005) The economics of service upgrades. J Serv Res 7(3):234–244. https://doi.org/10.1177/1094670504271148

Bront JJM, Méndez-Díaz I, Vulcano G (2009) A column generation algorithm for choice-based network revenue management. Oper Res 57(3):769–784. https://doi.org/10.1287/opre.1080.0567

Çakanyıldırım M, Özer Ö, Zhang X (2020) Dynamic pricing and timing of upgrades. Available at. https://doi.org/10.2139/ssrn.3056060

Chiang WC, Chen JC, Xu X (2007) An overview of research on revenue management: current issues and future research. Int J Rev Manag 1(1):97–128. https://doi.org/10.1504/IJRM.2007.011196

Cui Y, Duenyas I, Sahin O (2018) Pricing of conditional upgrades in the presence of strategic consumers. Manag Sci 64(7):3208–3226. https://doi.org/10.1287/mnsc.2017.2783

Cui Y, Orhun AY, Duenyas I (2019) How price dispersion changes when upgrades are introduced: theory and empirical evidence from the airline industry. Manag Sci 65(8):3835–3852. https://doi.org/10.1287/mnsc.2018.3117

Denizci Guillet B (2020) Online upselling: moving beyond offline upselling in the hotel industry. Int J Hosp Manag 84:102322. https://doi.org/10.1016/j.ijhm.2019.102322

Denizci Guillet B, Mattila AS, Peng ZL et al (2022) The joint influence of the timing and framing of an online upselling message on consumer perceptions: the roles of construal level and reactance. Int J Contemp Hosp Manag 34(7):2516–2534. https://doi.org/10.1108/ijchm-11-2021-1375
Dong L, Kouvelis P, Tian Z (2009) Dynamic pricing and inventory control of substitute products. Manuf Serv Oper Manag 11(2):317–339. https://doi.org/10.1287/msom.1080.0221
Gallego G, Stefanescu C (2009) Upgrades, upsells and pricing in revenue management. SSRN Electron J. https://doi.org/10.2139/ssrn.1334341
Gallego G, Iyengar G, Phillips R et al (2004) Managing flexible products on a network. CORC technical report TR-2004-01. Department of Industrial Engineering and Operations Research, Columbia University, New York
Gallego G, Li L, Ratliff R (2009) Choice-based EMSR methods for single-leg revenue management with demand dependencies. J Rev Pricing Manag 8(2–3):207–240. https://doi.org/10.1057/rpm.2008.53
Gallego G, Ratliff R, Shebalov S (2015) A general attraction model and sales-based linear program for network revenue management under customer choice. Oper Res 63(1):212–232. https://doi.org/10.1287/opre.2014.1328
Gönsch J, Steinhardt C (2015) On the incorporation of upgrades into airline network revenue management. Rev Manag Sci 9(4):635–660. https://doi.org/10.1007/s11846-014-0129-5
Klein R, Koch S, Steinhardt C et al (2020) A review of revenue management: recent generalizations and advances in industry applications. Eur J Oper Res 284(2):397–412. https://doi.org/10.1016/j.ejor.2019.06.034
Li Y, Rasekh L, Wang XB et al (2016) Stochastic model for hotel room pricing and upgrading. J Rev Pricing Manag 15(6):500–508. https://doi.org/10.1057/s41272-016-0066-9
Li Y, Mahmoudzadeh A, Wang XB (2022) Airlines seat pricing with seat upgrading. Multimodal Transp 1(4):100054. https://doi.org/10.1016/j.multra.2022.100054
Liu Q, Van Ryzin G (2008) On the choice-based linear programming model for network revenue management. Manuf Serv Oper Manag 10(2):288–310. https://doi.org/10.1287/msom.1070.0169
Mayer VF, Santos GEdO, Marques ORB (2022) Option framing for upselling tourism services: Does cognitive availability prevent irrational choices? Tour Econ 28(2):476–494. https://doi.org/10.1177/1354816620949722
McFadden D (1987) Regression-based specification tests for the multinomial logit model. J Econom 34(1–2):63–82. https://doi.org/10.1016/0304-4076(87)90067-4
Norvell T, Kumar P, Contractor S (2018) Assessing the customer-based impact of up-selling versus down-selling. Cornell Hosp Q 59(3):215–227. https://doi.org/10.1177/1938965518762836
Park CH, Yoon TJ (2022) The dark side of up-selling promotions: evidence from an analysis of cross-brand purchase behavior. J Retail 98(4):647–666. https://doi.org/10.1016/j.jretai.2022.03.005
Rusmevichientong P, Topaloglu H (2012) Robust assortment optimization in revenue management under the multinomial logit choice model. Oper Res 60(4):865–882. https://doi.org/10.1287/opre.1120.1063
Schaefers T, Leban M, Vogt F (2022) On-demand features: consumer reactions to tangibility and pricing structure. J Bus Res 139:751–761. https://doi.org/10.1016/j.jbusres.2021.09.054
Seng Pun C, Klabjan D, Karaesmen F et al (2016) Itinerary-based nesting control with upsell. J Rev Pricing Manag 15(2):107–137. https://doi.org/10.1057/rpm.2015.21
Shao L (2021) The pricing decision and channel choice of upgraded products in the presence of strategic consumers. Math Probl Eng 2021:1–10. https://doi.org/10.1155/2021/6674676
Steinhardt C, Gönsch J (2012) Integrated revenue management approaches for capacity control with planned upgrades. Eur J Oper Res 223(2):380–391. https://doi.org/10.1016/j.ejor.2012.05.047
Strauss AK, Klein R, Steinhardt C (2018) A review of choice-based revenue management: theory and methods. Eur J Oper Res 271(2):375–387. https://doi.org/10.1016/j.ejor.2018.01.011
Talluri K, Van Ryzin G (2004a) Revenue management under a general discrete choice model of consumer behavior. Manag Sci 50(1):15–33. https://doi.org/10.1287/mnsc.1030.0147
Talluri KT, Van Ryzin G (2004b) The theory and practice of revenue management. Springer, Berlin
Thirumuruganathan S, Al Emadi N, Jung S et al (2023) Will they take this offer? A machine learning price elasticity model for predicting upselling acceptance of premium airline seating. Inf Manag 60(3):103759. https://doi.org/10.1016/j.im.2023.103759
Train KE (2009) Discrete choice methods with simulation. Cambridge University Press, Cambridge
Wang X, Wang H, Zhang X (2016) Stochastic seat allocation models for passenger rail transportation under customer choice. Transp Res Part E Logist Transp Rev 96:95–112. https://doi.org/10.1016/j.tre.2016.10.003

Weatherford LR, Bodily SE (1992) A taxonomy and research overview of perishable-asset revenue management: yield management, overbooking, and pricing. Oper Res 40(5):831–844. https://doi.org/10.1287/opre.40.5.831

Weatherford LR, Ratliff RM (2010) Review of revenue management methods with dependent demands. J Rev Pricing Manag 9(4):326–340. https://doi.org/10.1057/rpm.2010.14

Yilmaz O, Pekgün P, Ferguson M (2017) Would you like to upgrade to a premium room? Evaluating the benefit of offering standby upgrades. Manuf Serv Oper Manag 19(1):1–18. https://doi.org/10.1287/msom.2016.0596

Yilmaz O, Ferguson M, Pekgün P et al (2022) Strategic behavior for hotel standby upgrade programs: empirical evidence and pricing implications. J Oper Manag 68(6–7):675–701. https://doi.org/10.1002/joom.1177

Zhang D, Adelman D (2009) An approximate dynamic programming approach to network revenue management with customer choice. Transp Sci 43(3):381–394. https://doi.org/10.1287/trsc.1090.0262

Davina Rauhaus received the B.Sc. degree in Industrial Engineering from the University of Duisburg-Essen in 2017, followed by an M.Sc. in 2018 from the same institution. She is currently pursuing a Ph.D. in Business Administration and Operations Research at the Mercator School of Management, University of Duisburg-Essen. Her research focuses on pricing strategies for upsells and ancillary services in the field of revenue management.

Jochen Gönsch obtained his Ph.D. in Operations Research from the University of Augsburg, Germany, in 2010. After his Ph.D., he continued working as a post-doctoral researcher at the Department of Analytics & Optimization. Since 2015, he is full professor for Service Operations at the Mercator School of Management at the University of Duisburg-Essen, Germany. His research focuses on pricing and revenue management, shared mobility, renewable energies and remanufacturing.

Claudius Steinhardt received his Ph.D. degree in Operations Research from the University of Augsburg, Germany. Since 2014, he is full professor for Business Analytics & Management Science at the University of the Bundeswehr Munich, Germany. He heads the working group "Pricing and Revenue Management" of the German Operations Research Society. In his current work, he focuses on demand management for last-mile logistics and shared mobility systems.

Authors and Affiliations

Davina Rauhaus[1] · Jochen Gönsch[1] · Claudius Steinhardt[2]

✉ Jochen Gönsch
jochen.goensch@uni-due.de

Davina Rauhaus
davina.hartmann@uni-due.de

Claudius Steinhardt
claudius.steinhardt@unibw.de

[1] Chair for Service Operations, Mercator School of Management, University Duisburg-Essen, Lotharstr. 65, 47057 Duisburg, Germany

[2] Chair of Business Analytics and Management Science, University of the Bundeswehr Munich, Werner-Heisenberg-Weg 39, 85579 Neubiberg, Germany