

Using flexible products to cope with demand uncertainty in revenue management

31.10.2009

Anita Petrick, Claudius Steinhardt, Jochen Gönsch, Robert Klein

Anita Petrick

*Chair of Operations Research, Institute of Business Administration,
Technische Universität Darmstadt, Hochschulstr. 1, 64289 Darmstadt*

e-mail: petrick@bwl.tu-darmstadt.de

Claudius Steinhardt (✉), Jochen Gönsch, Robert Klein

*Chair of Mathematical Methods, Institute of Statistics and Economic Theory,
Universität Augsburg, Universitätsstraße 16, 86159 Augsburg*

Phone: +49 821 598 4159

Fax: +49 821 598 4226

e-mail: claudius.steinhardt@wiwi.uni-augsburg.de

Jochen Gönsch

e-mail: jochen.goensch@wiwi.uni-augsburg.de

Robert Klein

e-mail: robert.klein@wiwi.uni-augsburg.de

Abstract

While flexible products have been popular for many years in practice, they have only recently gained attention in the academic literature on revenue management. When selling a flexible product, a firm retains the right to specify some of its details later. The relevant point in time is *after* the sale, but often *before* the provision of the product or service, depending on the customers' need to know the exact specification in advance. The resulting flexibility can help to increase revenues if capacity is fixed and the demand to come difficult to forecast. We present several revenue management models and control mechanisms incorporating this kind of flexible products. An extensive numerical study shows how the different approaches can mitigate the negative impact of demand forecast errors.

Keywords: Revenue Management, Flexible Products, Capacity Control

1 Introduction

Revenue management is often described as “a method which can help a firm to sell the right inventory unit to the right type of customer, at the right time and for the right price” (Kimes 2000). Revenue management approaches are based on theoretical foundations in areas like demand modeling, forecasting, and mathematical optimization. The resulting models and procedures are usually integrated into large revenue management systems, like automated booking systems, with the main purpose of restricting the availability of exactly specified products through some kind of capacity control. After accepting a request for a *specific product*, the selling firm is usually bound to its decision and has to reserve capacity in respect of the resources needed for production. If capacity is fixed and the demand forecast is imprecise, this may lead to a significant loss in revenue. In service industries, where the date of purchase and the date of provision are not identical, the firm can specifically gain additional flexibility by offering *flexible products*. Such products allow some of the details – especially those determining the resources necessary for production – to be specified after the sale when more demand has realized and there is less uncertainty. Nevertheless, from a customer’s point of view, it is often necessary, or at least desirable, to know the exact specification before the actual service provision. Consequently, the firm must give its customers advanced notice at a pre-agreed point in time between the sale and provision, which we call the *notification date*.

In practice, flexible products are offered, for instance, by tour operators or cruise lines. Having bought a flexible product, customers are only informed of their specific hotel, cabin, or even cruise itinerary shortly before the start of the journey. Furthermore, broadcasting companies often sell flexible products, allowing them to spontaneously schedule commercials. Other applications have been discussed in respect of air cargo or make-to-order environments, where the flexibility arises from given time windows that allow the supplier to autonomously arrange accepted requests in respect of time and resources.

The contribution of this paper is two-fold: On the one hand, we show how to integrate flexible products with an explicit notification date into existing capacity

control models and outline suitable extensions to common control strategies. On the other hand, we conduct an extensive simulation study, showing how the approaches presented can mitigate the negative impact of imprecise and biased forecasts by means of extended supply-side flexibility through flexible products.

The paper is organized as follows: In Section 2 we provide a brief introduction to the literature on revenue management in general and review the research on revenue management with any kind of supply-side flexibility in particular. Flexible products that allow advanced notice are formalized in Section 3. First, standard revenue management models are extended to include this type of flexible products. Second, we point out possible modifications to the control mechanisms applied during the booking process to take advantage of the flexibility. The results of our simulation study are discussed in Section 4. In Section 5, a summary and final conclusion are presented.

2 Literature review

There is an extensive literature on revenue management in general. For surveys, see, for example, Belobaba (1987), Weatherford and Bodily (1992), McGill and van Ryzin (1999), Tscheulin and Lindenmeier (2003) and Chiang et al. (2007). Furthermore, various textbooks provide an overview of the field as well as an in-depth discussion of certain aspects (see, e.g., Talluri and van Ryzin 2004; Phillips 2005). In the following, we focus on publications that are more related to our work.

Only few authors consider the application of revenue management techniques in areas where customers, in one way or another, are indifferent to the exact specification of products, thus allowing the firm to determine some details after the sale. One group of publications considers the case that the decision on the exact specification must be made immediately after the product has been sold; that is, the notification date equals the purchase date. Talluri (2001) considers an airline company serving several origin-destination (O&D) markets. He assumes O&D demand, implying that passengers are indifferent with regard to the multiple itineraries serving the same market as far as these are similar regarding attributes like arrival/departure time and price. A bid-price control is proposed, with customers

being assigned to an itinerary immediately after booking. Chen et al. (2003) investigate several approaches of O&D-based capacity control in air cargo revenue management, which can also be used for passenger transportation. Based on the opportunity costs of the associated resources in the network, alternative routes are generated dynamically. Special attention is paid to the adaption of bid-price controls to flexible products. Büke et al. (2008) present stochastic linear programming approximations for a setting with O&D demand and buy-ups.

However, if customers are somehow indifferent to different specification possibilities, there might be no need for the firm to sacrifice flexibility by determining the exact specification immediately after purchase. At best no advanced notice is needed at all, in which case notification can take place when the product is provided. Bartodziej and Derigs (2004) and Bartodziej et al. (2007) suggest approaches that assume that air cargo clients usually only ask for certain time-windows. During the booking process, it is consequently sufficient to ensure that all freight can be delivered in time. Client notification is not necessary, so that the booking requests are only definitely assigned to itineraries shortly before departure. The authors permit fractions of a booking to be delivered via different itineraries and use column generation to find cheap and feasible routings for O&Ds.

To the best of our knowledge, the only publications mentioning a notification date concept are Gallego and Phillips (2004) and Gallego et al. (2004). Among control approaches with notification at the end of the booking horizon, Gallego and Phillips (2004) also consider an approach that allows notification after the first of two periods. The authors explicitly introduce the concept of flexible products as a mostly inferior extension of higher valued specific products and outline possible applications ranging from air cargo to tour operators and multiple property management. Their algorithms compute booking limits by taking advantage of a special structure: The only flexible product offered guarantees the provision of one of two specific ones. Numerical studies identify increasing overall demand and improved capacity utilization as major advantages at the risk of cannibalizing higher valued demand. Gallego et al. (2004) consider the more general setting of a network with an arbitrary number of products. They develop models and algorithms for network revenue management problems under the usual assumption of inde-

pendent demand as well as the more general hypothesis of demand following a customer choice model. Additional flexible products with notification dates earlier than the end of the horizon are briefly discussed.

The two aforementioned publications consider a problem setting quite similar to ours. However, Gallego and Phillips (2004) analyze EMSR-based algorithms for two period problems from a rather theoretical point of view, while our approach has a more operational perspective and allows for an arbitrary number of periods. Compared to Gallego et al. (2004), we present a dynamic programming formulation including an arbitrary notification date within the booking horizon. Furthermore, the resulting certainty equivalent approximation leads to a linear program which differs from the one proposed by Gallego et al. (2004), as it explicitly accounts for the capacity requirements of requests for flexible products that are already accepted, thus allowing for their rearrangement.

The literature on upgrades and research on flexibility in make-to-order (MTO) environments are, in one way or the other, also related to our work. Upgrades are a very simple way of gaining flexibility, albeit by offering only specific products. Airlines, for example, sell tickets for specific flights and in case they need flexibility, they transport their passengers in a different – higher valued – compartment. If this is not possible and customers must be transported in a lower compartment (downgrade), or on a later flight, compensation has to be offered because this is done without explicitly informing the clients about this possibility at the time of purchase. Alstrup et al. (1986) have already considered the impact of up- and downgrades between two compartments on a single flight leg on revenue management. While some models only allow upgrades to the next higher product (see Netessine et al. 2002; Shumsky and Zhang 2009), others (see Karaesmen 2001; Karaesmen and van Ryzin 2004) are more general. Geraghty and Johnson (1997) investigate planned upgrades for a car rental company. In MTO environments, flexibility can stem from loose due dates. Gallien et al. (2004), for instance, base the acceptance of incoming orders on opportunity cost. Since they allow preemption, orders are simply scheduled according to their due date. Jalora (2006) focuses on scheduling and develops a heuristic to schedule orders in the period with the lowest opportunity cost. To decide on acceptance of orders in the

iron and steel industry, Spengler et al. (2007) take into account that products can be stored at different stages of the value chain. Kimms and Müller-Bungart (2007a) address a different application of flexible products, namely broadcasting companies' revenue management of their commercials, in which flexibility arises from the customers' indifference to various scheduling options.

The literature on decision postponement in supply chain management is also somewhat related to the idea of flexible products. Decision postponement aims at differentiating products as late as possible to allow companies to better cope with, among others, market shifts (see, e.g., Stadtler 2008, Chap. 1.3.1).

3 Revenue management with flexible products

In this section, we extend several well-known model formulations for network revenue management to consider flexible products with explicit notification dates as described in Section 1. Furthermore, we show how the proposed models are used within a capacity control mechanism over time. We provide a rather general description which does not emphasize special industry profiles, thereby underlining the scope of the presented concepts' applicability.

3.1 Setting and notation

We consider a firm disposing of a network $\mathcal{H} = \{1, \dots, l\}$ which consists of l resources with total capacities $\mathbf{C} = (C_1, \dots, C_l)$. Similar to the setting considered by Gallego et al. (2004) and Gallego and Phillips (2004), we assume a set of *specific products* $\mathcal{I} = \{1, \dots, n^s\}$ defined on \mathcal{H} , with the capacity consumption of a single unit of product $i \in \mathcal{I}$ expressed by the vector $\mathbf{a}_i = (a_{i1}, \dots, a_{il})$. The contribution margin r_i resulting from selling one unit of product i is given by the difference of revenue and variable costs. Furthermore, we consider a set of *flexible products* $\mathcal{J} = \{1, \dots, n^f\}$. For each $j \in \mathcal{J}$ there is a set $\mathcal{M}_j \subseteq \mathcal{I}$ describing its possible execution modes, which we assume to be a subset of the existing specific products without loss of generality. For flexible products, the contribution margin is again the difference between revenue and variable costs. In contrast to specific products, however, the variable cost v_{jm} of a flexible product j depends on the execution

mode $m \in \mathcal{M}_j$, and is therefore not uniquely defined. Thus, the contribution margin f_{jm} of a flexible product with revenue f_j is now also mode-specific.

Both product types must be sold within a given horizon, which can be sufficiently discretized into T time periods such that there is at most one buying request in each period t . The periods are numbered backwards in time, beginning with T , thereby assuming that any capacity remaining after period 1 is worthless.

As pointed out in Section 1, we focus on a generic case in which allocating specific resources to flexible product requests need not be done directly after the sale, nor only at the end of the booking horizon, but can take place any time in between. Therefore, we define one of the time periods as the notification date τ , during which all accepted flexible requests have to be assigned to a specific execution mode.

Hence, the total number of time periods can be separated according to the following three phases:

Within the periods $T, \dots, \tau + 1$, specific products as well as flexible ones may be sold. The probabilities of an arrival of a request for a specific product i and for a flexible product j in period t within this phase are denoted by $p_i(t)$ and $q_j(t)$, respectively. Consequently, the probability $p_0(t)$ of there being no incoming request in period t is expressed by $1 - \sum_{i \in \mathcal{I}} p_i(t) - \sum_{j \in \mathcal{J}} q_j(t)$ for $t \in \{T, \dots, \tau + 1\}$. The subsequent phase is the allocation period τ without any arrivals, during which all accepted flexible requests have to be assigned to a specific execution mode. In the third phase defined by the time periods $\tau - 1, \dots, 1$, no more flexible products may be sold. The probabilities for incoming specific requests i within this final phase are again denoted by $p_i(t)$, with the probability for no incoming request being defined by $1 - \sum_{i \in \mathcal{I}} p_i(t)$.

With each incoming request, the firm's problem is to decide whether to accept or reject it in order to maximize the expected overall contribution margin.

3.2 Dynamic programming model

We first develop a stochastic dynamic programming model for the firm's maximization problem, which leads to an *optimal policy* under the given model parameters. Therefore, we define the value function $V(t, \mathbf{c}, \mathbf{Y}^a)$ as the maximum ex-

pected contribution margin-to-go, which is dependent on the current state of the system $(t, \mathbf{c}, \mathbf{Y}^a)$. The state $(t, \mathbf{c}, \mathbf{Y}^a)$ is uniquely specified by the remaining periods t , the vector of the available resource capacities $\mathbf{c} = (c_1, \dots, c_l)$, and the vector $\mathbf{Y}^a = (Y_1^a, \dots, Y_{n_f}^a)$, which contains the total amount Y_j^a of requests that have already been accepted but not yet assigned to a specific execution mode for each flexible product $j \in \mathcal{J}$. Let \mathcal{Z} be the system's set of feasible states, where a state $(t, \mathbf{c}, \mathbf{Y}^a)$ is defined feasible if the remaining capacities \mathbf{c} are nonnegative and sufficient to satisfy all accepted flexible requests (see Gallego et al. 2004; Karaesmen and van Ryzin 2004).

Assuming independent demand, the value $V(t, \mathbf{c}, \mathbf{Y}^a)$ can now be computed recursively via the Bellman equation

$$V(t, \mathbf{c}, \mathbf{Y}^a) = \begin{cases} \begin{cases} \sum_{i \in \mathcal{I}} p_i(t) \cdot \max\{V(t-1, \mathbf{c}, \mathbf{Y}^a), r_i + V(t-1, \mathbf{c} - \mathbf{a}_i, \mathbf{Y}^a)\} \\ + \sum_{j \in \mathcal{J}} q_j(t) \cdot \max\{V(t-1, \mathbf{c}, \mathbf{Y}^a), f_j + V(t-1, \mathbf{c}, \mathbf{Y}^a + \mathbf{e}_j)\} \\ + p_0(t) \cdot V(t-1, \mathbf{c}, \mathbf{Y}^a) \end{cases} & \begin{matrix} \text{for } t = T, \dots, \tau + 1, \\ (t, \mathbf{c}, \mathbf{Y}^a) \in \mathcal{Z} \end{matrix} \\ \max_{\mathbf{y}^a} \left\{ -\sum_{j \in \mathcal{J}} \sum_{m \in \mathcal{M}_j} v_{jm} \cdot y_{jm}^a + V\left(t-1, \mathbf{c} - \sum_{j \in \mathcal{J}} \sum_{m \in \mathcal{M}_j} y_{jm}^a \cdot \mathbf{a}_m, \mathbf{0}\right) \right\} & \begin{matrix} \text{for } t = \tau, \\ (t, \mathbf{c}, \mathbf{Y}^a) \in \mathcal{Z} \end{matrix} \\ \left. \sum_{m \in \mathcal{M}_j} y_{jm}^a = Y_j^a \quad \forall j \in \mathcal{J}, y_{jm}^a \geq 0 \text{ and integer } \forall j \in \mathcal{J}, m \in \mathcal{M}_j \right\} & \\ \begin{cases} \sum_{i \in \mathcal{I}} p_i(t) \cdot \max\{V(t-1, \mathbf{c}, \mathbf{0}), r_i + V(t-1, \mathbf{c} - \mathbf{a}_i, \mathbf{0})\} \\ + p_0(t) \cdot V(t-1, \mathbf{c}, \mathbf{0}) \end{cases} & \begin{matrix} \text{for } t = \tau - 1, \dots, 1, \\ (t, \mathbf{c}, \mathbf{Y}^a) \in \mathcal{Z} \end{matrix} \end{cases} \quad (3.1)$$

with the boundary conditions

$$V(t, \mathbf{c}, \mathbf{Y}^a) = -\infty \quad \text{if } (t, \mathbf{c}, \mathbf{Y}^a) \notin \mathcal{Z}, \quad (3.2)$$

$$V(0, \mathbf{c}, \mathbf{Y}^a) = 0 \quad \text{if } (0, \mathbf{c}, \mathbf{Y}^a) \in \mathcal{Z}. \quad (3.3)$$

The formulation can be explained as follows: Within the interval defined by the periods $t = T, \dots, \tau + 1$, specific products as well as flexible ones may be sold. If a specific request i is accepted, the corresponding resources' capacities are immediately reduced, leading to the new state $(t-1, \mathbf{c} - \mathbf{a}_i, \mathbf{Y}^a)$. If a flexible request j is accepted, however, there is no direct capacity reduction, but the request has to be

memorized for later assignment. Thus, the next state is $(t-1, \mathbf{c}, \mathbf{Y}^a + \mathbf{e}_j)$, where \mathbf{e}_j denotes the j -th standard basis vector in \mathbb{R}^{n_f} . As in dynamic programming approaches to standard network revenue management problems, it is obvious that in an optimal control policy, a request for a specific product i is accepted if and only if its contribution margin exceeds its opportunity cost:

$$r_i \geq V(t-1, \mathbf{c}, \mathbf{Y}^a) - V(t-1, \mathbf{c} - \mathbf{a}_i, \mathbf{Y}^a). \quad (3.4)$$

A fairly similar result can be derived in respect of flexible products, namely that a request for j is accepted if and only if the following condition holds:

$$f_j \geq V(t-1, \mathbf{c}, \mathbf{Y}^a) - V(t-1, \mathbf{c}, \mathbf{Y}^a + \mathbf{e}_j). \quad (3.5)$$

Note that in this case, the term on the left-hand side, f_j , is simply the per-unit revenue, while the term on the right-hand side is the opportunity cost, which is implicitly based on the (expected) variable cost depending on the resource allocation taking place later in time at $t = \tau$.

In period $t = \tau$, the accepted flexible requests \mathbf{Y}^a are finally assigned to execution modes, where the variables y_{jm}^a quantify a feasible assignment of the accepted requests for each flexible product j to modes $m \in \mathcal{M}_j$. To shorten the notation, the variables y_{jm}^a are grouped in the vector \mathbf{y}_j^a for each $j \in \mathcal{J}$, with all the resulting vectors in turn being referenced by $\mathbf{y}^a = (\mathbf{y}_1^a, \dots, \mathbf{y}_n^a)$. The total variable cost resulting from a certain assignment is calculated as $\sum_{j \in \mathcal{J}} \sum_{m \in \mathcal{M}_j} v_{jm} \cdot y_{jm}^a$, with the next state being

$$\left(\tau - 1, \mathbf{c} - \sum_{j \in \mathcal{J}} \sum_{m \in \mathcal{M}_j} y_{jm}^a \cdot \mathbf{a}_m, \mathbf{0} \right).$$

In the following periods $t = \tau - 1, \dots, 0$, no flexible products are available for purchase, so that the resulting Bellman equation corresponds to standard formulations (see, e.g., Talluri and van Ryzin 1998; Bertsimas and Popescu 2003).

3.3 Linear approximations

Although of theoretical interest, the dynamic programming model presented in the previous section can barely be used to handle practical network sizes, due to the curse of dimensionality arising from the multidimensional state space. We therefore make use of a technique called certainty equivalent control (CEC), which belongs to the class of approximate dynamic programming methods (see Bertsekas 2005) and has also been successfully applied to standard capacity control (see Bertsimas and Popescu 2003). The realization of CEC in the context of our model (3.1)–(3.3) is straightforward and can be described as follows:

In the first step, an approximation is constructed for the value function $V(t, \mathbf{c}, \mathbf{Y}^a)$ by replacing all of the system's stochastic influences, namely the requests for each product arriving in the subsequent periods, with their expectations, and aggregating them over time. This is a very common approach in revenue management, leading to the well-known deterministic linear programming (DLP) model (see, e.g., Simpson 1989; Williamson 1992). In our setting with specific and flexible products, we have to take the expected aggregated demand-to-come \bar{D}_{it}^s and \bar{D}_{jt}^f , respectively, into consideration. As in the standard DLP formulation, we use decision variables x_i which – in case of $a_{hi} \in \{0,1\}$ for all $h \in \mathcal{H}$, $i \in \mathcal{I}$ – directly correspond to the number of units of capacity to be allocated to future requests for product i in respect of each required resource. The corresponding vector is $\mathbf{x} = (x_1, \dots, x_{n^s})$. In addition, we introduce decision variables y_{jm}^a and y_{jm}^e , with y_{jm}^a denoting capacity allocations to the accepted flexible requests in the different modes and y_{jm}^e denoting allocations to the expected future flexible requests. Again, the allocations are grouped in vectors \mathbf{y}^a and \mathbf{y}^e , respectively. The optimization problem can then be formulated as an extension of the standard DLP as follows (**DLP-ext**):

$$\tilde{V}^{DLP}(t, \mathbf{c}, \mathbf{Y}^a) = \max_{\mathbf{x}, \mathbf{y}^e, \mathbf{y}^a} \left(\sum_{i \in \mathcal{I}} r_i \cdot x_i + \sum_{j \in \mathcal{J}} \sum_{m \in \mathcal{M}_j} f_{jm} \cdot (y_{jm}^a + y_{jm}^e) \right) \quad (3.6)$$

subject to

$$\sum_{i \in \mathcal{I}} a_{hi} \cdot x_i + \sum_{j \in \mathcal{J}} \sum_{m \in \mathcal{M}_j} a_{hm} \cdot (y_{jm}^a + y_{jm}^e) \leq c_h \quad \text{for all } h \in \mathcal{H}, \quad (3.7)$$

$$\sum_{m \in \mathcal{M}_j} y_{jm}^a = Y_j^a \quad \text{for all } j \in \mathcal{J}, \quad (3.8)$$

$$\sum_{m \in \mathcal{M}_j} y_{jm}^e \leq \bar{D}_{jt}^f \quad \text{for all } j \in \mathcal{J}, \quad (3.9)$$

$$x_i \leq \bar{D}_{it}^s \quad \text{for all } i \in \mathcal{I}, \quad (3.10)$$

$$y_{jm}^a, y_{jm}^e \geq 0 \quad \text{for all } j \in \mathcal{J}, m \in \mathcal{M}_j, \quad (3.11)$$

$$x_i \geq 0 \quad \text{for all } i \in \mathcal{I}. \quad (3.12)$$

The objective function (3.6) maximizes the total contribution margin. The constraints (3.7) ensure that the remaining capacity \mathbf{c} is sufficient to produce allocations x_i , y_{jm}^a , and y_{jm}^e . The requirement that all accepted flexible requests Y_j^a for a particular product j are being allocated execution modes is met by constraints (3.8). Furthermore, the allocations x_i and y_{jm}^e should not exceed the expected demands (constraints (3.9); (3.10)) and all allocations must be nonnegative (constraints (3.11); (3.12)). Additionally, we define $\tilde{V}^{DLP}(t, \mathbf{c}, \mathbf{Y}^a) = -\infty$ if DLP-ext has no valid solution.

In the second step of the CEC approach, we simply use the values \tilde{V}^{DLP} obtained from the linear approximation within the decision rules (3.4) and (3.5), leading to the modified acceptance criteria

$$r_i \geq \tilde{V}^{DLP}(t-1, \mathbf{c}, \mathbf{Y}^a) - \tilde{V}^{DLP}(t-1, \mathbf{c} - \mathbf{a}_i, \mathbf{Y}^a) \quad (3.13)$$

and

$$f_j \geq \tilde{V}^{DLP}(t-1, \mathbf{c}, \mathbf{Y}^a) - \tilde{V}^{DLP}(t-1, \mathbf{c}, \mathbf{Y}^a + \mathbf{e}_j), \quad (3.14)$$

respectively.

Note that at the notification date τ , the optimization of model (3.6)–(3.12) is followed by the final assignment of all accepted flexible requests \mathbf{Y}^a . Therefore, the resource capacities are reduced according to the allocations y_{jm}^a delivered by the model. For $t < \tau$, we set $\bar{D}_{jt}^f = 0$ and $Y_j^a = 0$ for all $j \in \mathcal{J}$, so that the original formulation can be used again.

In an approach called randomized linear programming (RLP), Talluri and van Ryzin (1999) use the standard DLP with a Monte Carlo simulation in order to better incorporate the uncertainty of demand. To apply their approach in our setting (**RLP-ext**), we simply replace the standard DLP with the extended DLP formulation for flexible products (DLP-ext). Hence, we define the random variables D_{it}^s and D_{jt}^f , which refer to the aggregated demand-to-come for specific and flexible products, respectively. For each product, we generate K samples from the corresponding demand distribution, denoting them $D_{it}^{s_1}, \dots, D_{it}^{s_K}$ for all $i \in \mathcal{I}$ and $D_{jt}^{f_1}, \dots, D_{jt}^{f_K}$ for all $j \in \mathcal{J}$. For each $k = 1, \dots, K$, we then calculate an approximation of the value function \tilde{V}^{RLP_k} by solving an instance of model (3.6)–(3.12), replacing the expectations \bar{D}_{jt}^f and \bar{D}_{it}^s on the right-hand sides of constraints (3.9) and (3.10) with the sampled values $D_{jt}^{f_k}$ and $D_{it}^{s_k}$. Finally, an overall approximation of the value function is obtained by taking the average of the K values of the objective function:

$$\tilde{V}^{RLP}(t, \mathbf{c}, \mathbf{Y}^a) = \frac{1}{K} \sum_{k=1}^K \tilde{V}^{RLP_k}(t, \mathbf{c}, \mathbf{Y}^a). \quad (3.15)$$

3.4 Control mechanisms

For practical purposes, even the CEC approach is often too complex, as two instances of model (3.6)–(3.12) have to be solved in real-time for each incoming request. One possibility to simplify the control process is the usage of *bid-prices* – an approach initially proposed by Smith and Penn (1988), Simpson (1989), and Williamson (1992). This approach has evolved as one of the leading methods in practical applications due to its simplicity and relative robustness (see Talluri and van Ryzin 2004, Chap. 3.1.2.3). Bid-prices can be regarded as threshold prices in the sense that for each resource $h \in \mathcal{H}$, the corresponding bid-price π_h approximates the monetary value of a single capacity unit. Such bid-prices can be derived for deterministic model formulations by simply using the shadow prices of the corresponding capacity constraint (constraints (3.7) in our model) obtained from the optimal dual solution. In respect of model (3.6)–(3.12), it can then be

shown that a request for a specific product $i \in \mathcal{I}$ should be accepted if and only if the following conditions hold:

$$r_i \geq \sum_{h \in \mathcal{H}} a_{hi} \cdot \pi_h, \quad (3.16)$$

$$a_{hi} \leq c_h - \sum_{j \in \mathcal{J}} \sum_{m \in \mathcal{M}_j} a_{hm} \cdot y_{jm}^a \quad \text{for all } h \in \mathcal{H}. \quad (3.17)$$

Note that compared to (3.13), the opportunity cost of a request is now approximated by the (quantity-adjusted) sum of the required resources' bid-prices (Condition (3.16)). Condition (3.17) guarantees that there is enough capacity left to fulfill the current request as well as all flexible requests that have already been accepted. If conditions (3.16) and (3.17) lead to an acceptance decision for product i , the corresponding resources are adjusted accordingly, namely $c_h := c_h - a_{hi}$ for all $h \in \mathcal{H}$. Although strongly dependent on the remaining capacity and demand-to-come, bid-prices are not usually recalculated for every request, but are kept constant for a certain period of time. The frequency of bid-price updates depends on the available data as well as on the computational effort, and is often adapted to typical demand patterns observed in the particular field of application. Note that bid-prices can also be consistently derived from the RLP by simply averaging the shadow prices obtained from the K different model instances (see Talluri and van Ryzin 1999).

The bid-price approach can be used to handle requests for flexible products as well. Therefore, we generalize the acceptance criteria as follows: An incoming request for a flexible product $j \in \mathcal{J}$ should be accepted if and only if there is at least one execution mode $m' \in \mathcal{M}_j$ with

$$f_{jm'} \geq \sum_{h \in \mathcal{H}} a_{hm'} \cdot \pi_h \quad (3.18)$$

and

$$a_{hm'} \leq c_h - \sum_{j \in \mathcal{J}} \sum_{m \in \mathcal{M}_j} a_{hm} \cdot y_{jm}^a \quad \text{for all } h \in \mathcal{H}. \quad (3.19)$$

If the request is accepted, it should be assigned to a mode $m^* \in \mathcal{M}_j$ with

$$m^* = \arg \max_{m' \in \mathcal{M}_j} \left\{ f_{jm'} - \sum_{h \in \mathcal{H}} a_{hm'} \cdot \pi_h \mid \text{subject to (3.19)} \right\}. \quad (3.20)$$

The consistency of this acceptance (and assignment) criteria with an optimal solution of model (3.6)–(3.12) can be shown by duality theory (see appendix A). The result is quite intuitive as it reflects the idea that a flexible request should be assigned so that the maximum benefit resulting from the mode-specific contribution margin $f_{jm'}$, on the one hand and the (approximated) opportunity cost $\sum_{h \in \mathcal{H}} a_{hm'} \cdot \pi_h$ on the other is achieved. If conditions (3.18)–(3.20) lead to an acceptance decision in respect of request j in mode m^* , the correspondent variable that stores the number of accepted flexible requests in mode m^* must be adjusted accordingly, namely $y_{jm^*}^a := y_{jm^*}^a + 1$.

As long as the notification date τ is not exceeded, the resource allocation in respect of flexible products is only temporary: Each time a bid-price update is performed by resolving model (3.6)–(3.12), all accepted requests for each flexible product $j \in \mathcal{J}$ are consolidated in the corresponding model parameter Y_j^a . Thus, they can be rearranged according to the current demand forecast and capacity utilization, leading to new (temporary) allocations y_{jm}^a .

In the following, we briefly outline two other control mechanisms which can easily be generalized to the flexible product setting. The first approach is based on *booking limits*, which, roughly speaking, denote the maximum number of requests that should be accepted in respect of each of the products. One straightforward way of deriving such limits is to directly use the future requests' allocations produced by DLP-ext. However, as this generally results in a large number of very small allocations, it is common to define a nesting structure based on the different products' resource-specific rankings in order to grant higher valued products access to capacity units originally intended for lower valued products. In order to obtain resource-specific ranking hierarchies, the total product revenue is usually split between the resources, for example, according to mileage, total number of needed resources, or the relative revenue value of local demand for each resource.

A resource-specific booking limit b_{hi} for product i on resource h is then derived by subtracting all the allocations in respect of higher valued products from the remaining capacity c_h . During the booking process, the control is completely based on these pre-calculated booking limits in the sense that a request for product i is accepted if $b_{hi} \geq a_{hi}$ for all $h \in \mathcal{H}$. After acceptance, the corresponding booking limits must be reduced accordingly.

Flexible products can be easily incorporated into this framework by using the mode-specific allocations y_{jm}^e obtained from model (3.6)–(3.12). The different execution modes are then considered separately within the rankings, leading to mode-specific booking limits b_{hjm} . Consequently, a request for product $j \in \mathcal{J}$ should be accepted if there is a mode $m' \in \mathcal{M}_j$ with $b_{hjm'} \geq a_{hm'}$ for all $h \in \mathcal{H}$. Accepted requests for flexible products are temporarily assigned to a mode with a maximum positive booking limit and – like the bid-price approach described before – can be rearranged as soon as model (3.6)–(3.12) is reoptimized.

The second approach has recently been proposed by Topaloglu (2009), who in a way combines the idea of booking limits with the simplicity of a bid-price control. Basically, a standard bid-price approach is performed, with the exception that if equality holds in condition (3.16), the current specific request is only accepted with probability x_i / \bar{D}_{it}^s . The intuition behind this is that in case of equality, the allocation in respect of product i can be smaller than the expected demand-to-come, therefore a pure booking limit control would not accept all incoming requests. Consequently, a flexible product request j satisfying conditions (3.18) and (3.19) should only be accepted with probability $\sum_{m \in \mathcal{M}_j} y_{jm}^e / \bar{D}_{jt}^f$ if

$$f_{jm^*} = \sum_{h \in \mathcal{H}} a_{hm^*} \cdot \pi_h,$$

with m^* being determined by (3.20).

4 Simulation study

In this section, we analyze the performance of the proposed controls based on the results of an extensive numerical study. All implementations were undertaken by means of “Java 2 Platform, Standard Edition” version 1.6.0 by Sun Microsystems.

The simulations ran on a PC with two 3 GHz Intel Pentium processors, 1 GB RAM and Windows XP.

4.1 Simulation environment

To evaluate the control mechanisms, we introduce several problem classes for which different instances are automatically generated. Each simulation run comprises a complete booking period with requests for all products. The problem classes considered are presented in the context of passenger airline revenue management, but we believe that the two underlying resource networks also occur in a wide range of other areas.

- *Network 1* consists of independent flight legs without the possibility of combining two flights and change aircraft. This is typical of low cost airlines offering point to point service. Examples of businesses with similar resource structures include tour operators with multiple hotels, cruise ships, or manufacturing processes with several production technologies.
- *Network 2* is more complex and represents a part of the network of a full service carrier. The distinctive property is that passengers can now combine flight legs to form connecting flights. Apart from aviation, comparable structures can arise, for example, in production if goods need processing on more than one machine.

In the following, the two networks are presented in more detail. We obtain eight problem classes by defining four demand scenarios for each network.

4.1.1 Network 1

Network 1 consists of four similar and independent flight legs (e.g., flights connecting two cities at different times of day). We assume that these legs have a capacity of $C_h = 200$ seats each. As no connections are possible in this simple network, customers can only choose between four itineraries, each consisting of exactly one flight leg. We consider only one compartment with four booking classes per flight leg, making up a total of 16 products. The resource consumption is the same for every booking class: A passenger necessitates $a_{hi} = 1$ seats on the leg he uses, otherwise $a_{hi} = 0$. On all itineraries, the products have an expected

demand ratio of 1:2:3:4 and are priced at 550, 400, 300, and 210, depending on the booking class. In addition to these specific products, a flexible product is offered that assures transportation on one of the flights. As the passenger has no influence on the flight the airline chooses for him, this product will most likely be regarded as inferior. As compensation, it is priced at 168; that is, 20% below the cheapest booking class. The demand for this flexible product is 15% of total demand.

Revenue management's potential benefits are highly dependent on the scarcity of resources. Therefore, for each leg h , a *nominal load factor* (LF_h) is defined as the total expected demand for seats divided by the total number of available seats. By specifying these values as input, different demand situations are represented and random demand is generated accordingly (see Kimms and Müller-Bungart 2007b; Klein 2007). For Network 1, we define four demand scenarios reflecting high (1-H), medium (1-M), low (1-L), and erratic (1-E) demand with an average nominal load factor of 1.4, 1.1, 0.9, and 1.0, respectively. The individual values of LF_h in respect of each scenario are given in appendix B.

4.1.2 Network 2

The second network is part of a traditional airline's network: Within a single day, 15 flight legs connect the cities A, B, C, and D (see Fig. 1). There are six short haul (1–6), seven medium haul (7–13), and two long haul flights (14, 15). The short haul flights go from A to B (1–3) and C to D (4–6). The medium haul flights connect A with C (7–9) and B with D (10–13). D can be reached from A directly with the two long haul flights (14, 15). In addition to 15 itineraries corresponding to one of these flight legs (continuous arcs in Fig. 1), nine itineraries, which consist of two legs – one short and one medium haul one – are available. The itineraries connect A and D, including a transfer in B or C (denoted by dashed arcs connecting two legs). Again, we consider only one compartment with four booking classes and demand ratios as in Network 1. Combining these booking classes with the 24 itineraries defines 96 specific products.

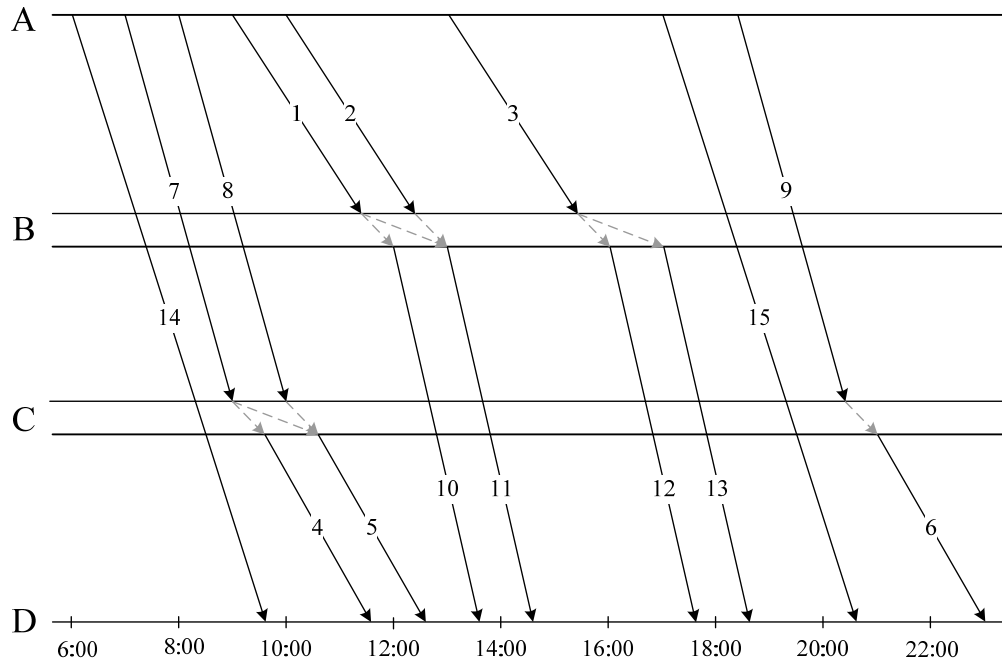


Fig. 1 Network 2 (time-space network)

The ticket prices depend on the length and type of itinerary. The most expensive are the direct long haul flights, priced at 1200, 950, 700, and 500, depending on the booking class. Furthermore, we assume the airline gives a discount to passengers travelling from A to D who accept the hassle of transferring. They save 8% compared to a direct flight from A to D. For a single medium or short haul flight, passengers are charged 80% (respectively 30%) of the corresponding booking class on a long haul flight.

In addition, very price sensitive customers can choose one of the following three flexible products connecting A to D at 25% below the fare for a direct flight in the cheapest booking class:

- The first flexible product “direct flight from A to D with flexible departure time” does not require the passengers to change aircraft but implies flying either early in the morning (14) or late in the evening (15).
- Arrival in the afternoon is guaranteed by “one-stop flight from A to D, early departure.” This product offers transportation by an itinerary consisting of the following pairs of flights: (1,10), (1,11), (7,4), or (7,5).
- The appropriate product for travelling in the evening is “one-stop flight from A to D, late departure” with transportation on (3,12), (3,13), or (9,6).

In all three cases, the demand for the flexible product equals 20% of the total demand for the respective itineraries.

Similar to Network 1, we define nominal load factors LF_h in respect of four demand situations: high (2-H), medium (2-M), low (2-L), and erratic (2-E) demand with capacity weighted average nominal load factors of 1.3, 1.1, 0.9, and 1.1, respectively. With the individual LF_h values for each demand situation, capacities C_h , and some additional specifications regarding the proportion of connecting passengers, the expected total demand \bar{D}_{iT}^s for each specific product (\bar{D}_{jT}^f for flexible products, respectively) is well defined. As Network 2 is quite complex, the respective values and calculation formulae are given in appendix C.

4.1.3 Demand & forecast

The most important step to obtain simulation runs is generating incoming booking requests for each test case. We describe the creation of requests in the context of specific products i . The procedure for flexible products j is analogous. Requests for each product arrive according to a nonhomogeneous Poisson process, a common assumption in revenue management (see Kimms and Müller-Bungart 2007b). The corresponding intensity functions $\lambda_i(t)$, called booking curves, have a triangular shape (see Klein 2007). Demand for product i occurs in the interval $[t_i^S, t_i^F]$. Beginning with t_i^S , the demand intensity rises linearly from $\lambda_i(t_i^S) = 0$ until it reaches the maximum value at t_i^{max} . Subsequently, it declines steadily until t_i^F with $\lambda_i(t_i^F) = 0$. Given total expected demand for product i as well as the values t_i^S , t_i^{max} , and t_i^F , the intensity function $\lambda_i(t)$ is well defined and requests can be generated efficiently using standard procedures (see, e.g., Law 2007, Chap. 8.6.2).

As flight tickets are usually available one year in advance, we set the length of the booking period to 360 days. To shorten the specification, identical values of t_i^S , t_i^{max} and, t_i^F are used for all products based on the same booking class. The choice of these parameters reflects that more expensive booking classes are generally sold near departure (see Table 1).

Table 1 Parameters of the booking curves

	Booking class				
	1	2	3	4	flexible
t_i^S	60	90	270	360	120
t_i^{max}	2	30	50	80	50
t_i^F	0	5	15	30	40

To incorporate forecast accuracy in our simulations, forecasting errors (see, e.g., Talluri and van Ryzin 2004, Chap. 9.5) can be controlled by an *itinerary-based bias* or a *systematic bias*, which distort the expectation of demand \bar{D}_{it}^s (\bar{D}_{jt}^f , respectively) used in the revenue management models after the requests have been generated as outlined above. The *itinerary-based bias* allows some itineraries to have higher demand than expected, while others are less popular than predicted. For this purpose, an upper bound δ_1 is specified. For each itinerary, a random value $\hat{\delta}_1 \sim U(-\delta_1, \delta_1)$ is drawn and the expected demand for every specific product i based on this itinerary is distorted by $\hat{\delta}_1$. That is, $(1 + \hat{\delta}_1) \cdot \bar{D}_{it}^s$ is used as a forecast of the expected aggregated demand-to-come. The expectation \bar{D}_{jt}^f is altered accordingly, but with independent $\hat{\delta}_1$ drawn for each flexible product j . The intuition behind the second variant, the *systematic bias*, is a general over- or underestimation of demand. Similar to Lee (1990) as well as to Weatherford and Bodily (1992), a deterministic value $\delta_2 \in [-0.5, 0.5]$ is used as input. Demand is then wrongly expected to be $(1 + \hat{\delta}_2) \cdot \bar{D}_{it}^s$ and $(1 + \hat{\delta}_2) \cdot \bar{D}_{jt}^f$ for specific and flexible products, respectively.

4.2 Results

In the following, we present the numerical results of an extensive simulation study to analyze the impact of flexible products on revenue management in uncertain environments in general, as well as to assess different control mechanisms' performance in particular. As the stochastic dynamic programming approach presented in Section 3 is computationally intractable even for small instances, we restrict ourselves to the analysis of corresponding approximations based on certainty equivalent control. Therefore, we implement five different control mechanisms $\mathcal{S} = \{\text{CEC}, \text{BL-DLP}, \text{BP-DLP}, \text{BP-RLP}, \text{T-DLP}\}$. While the first, CEC, is the direct application of certainty equivalent control – as described in Section 3.3

– using model ((3.6)–(3.12)) to approximate the value function, the other mechanisms further reduce the control effort (see Section 3.4): The second approach is an implementation of booking limits based on allocations obtained from the DLP-ext (BL-DLP). We follow de Boer et al. (2002) in respect of the nesting structure and base the ranking on the values

$$\bar{r}_{hi} = r_i - \sum_{h \in \mathcal{H}} a_{hi} \cdot \pi_h,$$

which can be regarded as an approximation of the resource-related net benefit of accepting a specific product i using resource h . Consequently, the net benefit of flexible product j in mode m is given by

$$\bar{f}_{hjm} = f_{jm} - \sum_{h \in \mathcal{H}} a_{hm} \cdot \pi_h.$$

Booking limits are adapted according to *standard nesting* (see, e.g., Talluri and van Ryzin 2004, Chap. 2.1.1.3). The third and fourth approaches are bid-price controls based on DLP-ext (BP-DLP) and RLP-ext (BP-RLP), respectively. In line with the literature (see, e.g., Talluri and van Ryzin 1999), we did some preliminary testing to determine the sample size for RLP-ext. It showed that the performance of this approach did not improve significantly with sample sizes larger than $K = 20$. The last approach is the implementation of Topaloglu’s acceptance criteria in which the underlying bid-prices are calculated with DLP-ext (T-DLP). Bid-prices as well as booking limits are (re)calculated 13 times during the booking horizon by reoptimizing the underlying model 360, 180, 110, 80, 65, 50, 40, 30, 20, 10, 5, 3, and 1 days before departure, thereby increasing the frequency towards the end of the booking horizon when demand intensity is higher. Unless otherwise noted, the notification date of all flexible products is 15 days before departure.

The simulation study is based on the eight problem classes introduced in Section 4.1. For each test case, 200 independent simulation runs are performed. We introduce several analytical measures to judge the performances. The first is the *optimality gap* (see, e.g., de Boer et al. 2002), which, for each simulation run

$\kappa = 1, \dots, 200$, measures the relative deviation of the value \tilde{V}_κ^s achieved by means of a certain control mechanism $s \in \mathcal{S}$ from the optimal value V_κ^{Opt} , which is hypothetically realized under complete demand information:

$$OG_\kappa^s = \frac{(\tilde{V}_\kappa^s - V_\kappa^{Opt})}{V_\kappa^{Opt}} \cdot 100\%. \quad (4.1)$$

V_κ^{Opt} is calculated by solving DLP-ext after observing the demand stream and thereby using the real demand realization instead of the expected values in (3.9) and (3.10).

To make a comparison of the results obtained from the different control mechanisms even easier, we additionally introduce the *performance gap*, which analogously measures the relative deviation of the values obtained from two different control mechanisms s and s' :

$$PG_\kappa^{s,s'} = \frac{(\tilde{V}_\kappa^s - \tilde{V}_\kappa^{s'})}{\tilde{V}_\kappa^{s'}} \cdot 100\%. \quad (4.2)$$

Furthermore, we define

$$CU_\kappa^s = \frac{1}{l} \cdot \sum_{h \in \mathcal{H}} R_{hk}^s \quad (4.3)$$

as the average *capacity utilization* resulting from a certain control mechanism s over all the resources in the considered network, with R_{hk}^s being the quotient of occupied and total leg capacity C_h .

To measure the performance of a complete simulation consisting of 200 runs, we compute the corresponding average values \overline{OG}^s , $\overline{PG}^{s,s'}$ and \overline{CU}^s .

4.2.1 Performance evaluation of controls for flexible products

We begin with a basic performance analysis of the five control mechanisms in respect of different forecast qualities. Table 2 shows the values of \overline{OG}^s for the eight problem classes described in Section 3.1, with forecast errors induced by the itinerary-based bias. The results show that as the upper bound δ_1 on the bias in-

creases from zero (very good forecast) to one (poor forecast), the optimality gap increases as well because all the controls suffer from a poor forecast. As expected, CEC clearly outperforms the other mechanisms in respect of all the test cases and values of δ_1 , with the exception of three, in which its performance is only slightly less good than that of the best. Nevertheless, CEC's computational effort which results from the online calculation of two value functions for every incoming request is prohibitive for real-world use. In our test cases, the runtime was about four to nine times that of BP-DLP. The BL-DLP results are also very good. Direct control through booking limits allows the number of accepted low fare requests to be accurately controlled, especially when demand is high.

Table 2 Comparison of the controls for different forecast qualities (\overline{OG}^s)

δ_1	CEC	BL-DLP	BP-DLP	BP-RLP	T-DLP	CEC	BL-DLP	BP-DLP	BP-RLP	T-DLP
	1-H					2-H				
0	-2.47	-2.63	-4.76	-3.22	-2.50	-1.93	-2.08	-3.21	-2.17	-2.04
0.25	-3.55	-3.93	-5.48	-4.50	-3.85	-2.80	-3.13	-3.84	-3.40	-3.26
0.5	-5.24	-5.84	-6.86	-6.50	-6.15	-4.48	-5.08	-5.56	-5.62	-5.52
0.75	-7.14	-8.07	-8.58	-9.04	-8.95	-6.44	-7.30	-7.70	-8.11	-8.02
1	-9.22	-10.56	-11.13	-11.66	-11.83	-8.44	-9.53	-9.86	-10.41	-10.53
	1-M					2-M				
0	-1.82	-2.02	-2.63	-2.10	-2.01	-1.53	-1.69	-2.34	-1.82	-1.69
0.25	-2.41	-2.82	-3.07	-3.02	-3.02	-2.10	-2.38	-2.74	-2.58	-2.63
0.5	-3.43	-4.09	-3.90	-4.48	-4.56	-3.38	-3.86	-3.94	-4.26	-4.37
0.75	-4.76	-5.71	-5.37	-6.35	-6.50	-4.47	-5.14	-5.10	-5.72	-5.89
1	-5.71	-7.59	-7.28	-8.34	-8.83	-5.94	-6.88	-6.62	-7.63	-7.93
	1-L					2-L				
0	-0.85	-1.25	-0.81	-0.79	-0.84	-0.89	-1.06	-1.04	-1.07	-1.02
0.25	-0.80	-1.18	-0.75	-0.91	-1.13	-1.14	-1.38	-1.29	-1.41	-1.55
0.5	-1.21	-1.86	-1.17	-1.61	-2.23	-1.89	-2.28	-2.09	-2.47	-2.75
0.75	-2.09	-3.25	-2.36	-3.44	-4.08	-2.59	-3.24	-2.84	-3.62	-3.99
1	-3.11	-5.30	-4.51	-6.03	-6.52	-3.43	-4.36	-3.89	-4.73	-5.26
	1-E					2-E				
0	-1.39	-1.70	-2.16	-1.63	-1.47	-1.60	-1.74	-2.90	-1.94	-1.70
0.25	-1.93	-2.34	-2.34	-2.41	-2.35	-2.34	-2.63	-3.54	-2.95	-2.71
0.5	-2.87	-3.39	-3.18	-3.66	-3.83	-3.70	-4.25	-4.68	-4.57	-4.49
0.75	-3.96	-4.66	-4.39	-5.31	-5.71	-5.32	-6.00	-6.31	-6.49	-6.67
1	-5.04	-6.27	-5.83	-7.68	-7.76	-6.87	-7.70	-7.90	-8.40	-8.75

However, nested booking limits exhibit various disadvantages in complex airline networks (see, e.g., Talluri and van Ryzin 2004, Chap. 3.1.2). Hence, in the following, we focus on the common bid-price based mechanisms and compare BP-DLP, BP-RLP, and T-DLP in more detail. To make the comparison more unders-

tandable, the best value obtained by one of these three approaches is printed in bold in Table 2. Given a good forecast, T-DLP obviously yields better results than BP-RLP, which in turn performs better than BP-DLP. However, T-DLP and BP-RLP are dependent on the forecast quality, whereas BP-DLP behaves more robustly. Hence, for each of the test cases considered, there is a certain threshold for δ_1 where this order is reversed: If the forecast is less accurate, BP-DLP outperforms the other two, with BP-RLP having slight advantages over T-DLP. The threshold depends on the relative demand: The higher the demand, the longer T-DLP is superior. Regardless of forecast errors, the T-DLP's capacity utilization is always below that of BP-DLP and BP-RLP. In respect of good forecasts, this underlines the well-known trade-off between high capacity utilization and revenue maximization. Furthermore, it shows that T-DLP succeeds in protecting capacity from low value requests and is thus able to accept more high value demand.

We now examine the performance of the controls in respect of forecasts with systematic bias δ_2 , instead of the itinerary-based bias. Table 3 shows the results of four test cases. In addition, they are exemplarily illustrated in respect of the 2-E network in Fig. 2. The fundamental differences between the over and underestimation of demand can be observed at a glance: When low demand is misleadingly expected due to a negative bias, the controls particularly tend to accept all requests until the capacity limit is reached, which is similar to a first-come-first-served approach. As this mostly leads to low value requests occupying too much capacity, approaches accepting fewer requests – CEC, T-DLP and BL-DLP – yield better results. However, if there is a positive bias, selling more low valued tickets is an advantage and BP-DLP performs well, sometimes even better than CEC. Generally, when demand is high, the controls are more sensitive to a negative bias, and when demand is low, they are more sensitive to overestimation.

Table 3 \overline{OG}^s with forecast bias

δ_2	CEC	BL-DLP	BP-DLP	BP-RLP	T-DLP	CEC	BL-DLP	BP-DLP	BP-RLP	T-DLP
	1-H					2-H				
-0.5	-14.13	-15.10	-17.84	-15.89	-15.08	-12.27	-13.23	-15.20	-13.81	-13.69
-0.25	-7.31	-7.80	-11.12	-8.06	-7.14	-6.55	-7.17	-9.13	-7.02	-6.85
0	-2.44	-2.61	-4.68	-3.24	-2.61	-1.94	-2.12	-3.23	-2.17	-2.05
0.25	-4.59	-5.46	-4.03	-7.44	-5.46	-4.82	-5.63	-4.69	-7.12	-7.32
0.5	-9.19	-10.69	-9.02	-13.03	-10.69	-10.32	-12.22	-10.77	-14.60	-15.03
	1-M					2-E				
-0.5	-6.21	-6.71	-7.23	-7.18	-6.86	-9.91	-10.64	-11.89	-10.94	-10.97
-0.25	-4.08	-4.42	-5.29	-4.56	-4.45	-5.43	-6.00	-7.83	-6.53	-5.80
0	-1.85	-2.05	-2.72	-2.17	-1.99	-1.61	-1.75	-2.94	-2.00	-1.69
0.25	-3.67	-4.52	-3.20	-5.56	-6.44	-4.14	-4.93	-3.98	-5.68	-6.20
0.5	-7.82	-9.79	-8.10	-11.48	-12.83	-8.68	-9.83	-8.99	-11.09	-13.06




Fig. 2 \overline{OG}^s for 2-E

To sum up, we note that CEC almost always yields the highest revenue. However, like BL-DLP, it is often not suitable for real-world application. BP-DLP is fast and quite robust in respect of forecast errors. However, in respect of good forecasts, it is outperformed by T-DLP and BP-RLP. If demand is systematically overestimated, BP-DLP leads to good results, while in case of underestimation, CEC, T-DLP, and BL-DLP achieve more revenue.

4.2.2 Performance gain of flexible products

In the following, we investigate the potential revenue benefits which arise from offering flexible products in detail. We first focus on the flexibility's revenue contribution in general. One obvious approach would be to compare settings with and

without flexible products. However, by doing so, the results obtained are strongly dependent on application-specific assumptions about consumer behavior. These include demand induction resulting from offering a cheaper (flexible) product and from “cannibalization”, which can be regarded as the propensity to buy a flexible product instead of a more expensive specific one. To exclude these marketing effects, we stick to the test cases used before, instead of creating new test cases without flexible products, and compare our controls with variants that merely do not make use of the flexibility. Through straightforward modifications, these variants, which we call ad-hoc (AH), irrevocably assign requests for flexible products to execution modes immediately after acceptance. We restrict the analysis to the controls CEC, BP-DLP, and BP-RLP and use the measure $\overline{PG}^{s,s'}$ to compare them to the corresponding ad-hoc controls, where s' is the ad-hoc variant of s .

The results are summarized in Table 4. It turns out that exploiting the flexibility yields an increase of up to 4% in revenue. Note that this figure is reliable at the 99% confidence level, assuming normal distribution of the mean $\overline{PG}^{s,s'}$ over the 200 simulation runs. In particular, for $\delta_1=1$, the confidence intervals are $[3.97; 5.32]$, $[3.99; 5.30]$ and $[4.12; 5.46]$ in respect of BP-DLP (1-H), BP-RLP (1-H) and CEC (1-M), respectively.

The exact revenue gain depends on forecast and demand. In particular, the lower the forecast quality, the more important it becomes to rectify the misallocations of flexible requests accepted early in the booking horizon and, thus, free up scarce resources for high value requests. BP-DLP yields the biggest gains in 34 of the 40 test cases, with BP-RLP enjoying only slightly lower gains. In respect of poor forecasts, even CEC, which by construction adapts better to forecast errors anyway, benefits considerably from the flexibility. In respect of good forecasts, the increase is much lower, as the expensive approximation of the opportunity cost results in comparatively good allocations for the ad-hoc variant of CEC. Fig. 3 visualizes these results for 1-H.

Table 4 Revenue increase due to flexibility in percent ($\overline{PG}^{s,s'}$)

δ_i	1-H			1-M			1-L			1-E		
	CEC	BP-DLP	BP-RLP	CEC	BP-DLP	BP-RLP	CEC	BP-DLP	BP-RLP	CEC	BP-DLP	BP-RLP
0	0.18	0.67	0.61	0.34	1.02	0.93	0.74	1.88	0.10	0.29	2.03	0.09
0.25	0.38	1.05	0.97	0.72	1.44	1.31	0.99	2.07	0.65	0.41	2.24	0.58
0.5	1.07	2.52	2.01	1.97	2.33	2.17	2.31	2.43	1.47	0.82	2.49	1.34
0.75	2.43	3.80	3.76	2.79	2.88	2.92	3.01	2.04	1.87	1.20	3.13	2.07
1	3.68	4.65	4.65	4.79	2.80	3.13	4.25	1.20	1.29	1.93	3.47	2.39
	2-H			2-M			2-L			2-E		
0	0.12	0.25	0.10	0.14	0.40	0.25	0.12	0.47	-0.02	0.07	0.23	0.21
0.25	0.28	0.53	0.27	0.33	0.68	0.47	0.27	0.62	0.20	0.14	0.38	0.37
0.5	0.62	1.07	0.70	0.79	1.14	1.02	0.85	0.95	0.67	0.40	0.76	0.69
0.75	1.12	1.56	1.32	1.25	1.54	1.45	1.06	1.13	0.98	0.72	1.03	1.02
1	1.60	1.86	1.75	1.63	1.73	1.70	1.37	0.95	1.01	1.01	1.36	1.34




Fig. 3 Ad-hoc vs. flexible for 1-H ($\overline{PG}^{s,s'}$)

To gain further insight, we now analyze the impact of the notification date on the revenue benefits obviously achieved by the flexibility. As pointed out in Section 1, the notification date is an essential part of a flexible product's definition, specifying the point in time when the firm has to determine the flexible requests' execution mode and inform the customers. While until now we have assumed that this decision happens 15 days before departure, in the following, ceteris paribus, the notification date is varied between the end of the flexible products' selling period (40 days before departure) and the end of the booking horizon. In line with our investigation of the ad-hoc variant, we assume demand to be independent of the notification date. Thus, in respect of each control, the simulation runs do not differ until 40 days before departure and the controls BP-DLP respectively come to the

same decision. At the notification date, however, the decision about the flexible requests' allocation has finally to be made. This decision is based on the forecast available at that time and no reallocations are possible afterwards. Thus, with a later notification date, the same number of flexible requests is accepted, but potentially more specific requests can be acknowledged if scarce resources remain available by flexible requests being moved, thereby leading to higher capacity utilization as well.

Table 5 shows the optimality gap as well as the capacity utilization for three different forecast qualities $\delta_1 \in \{0, 0.25, 0.5\}$ and the problem classes 1-M and 2-M. Apparently, a later notification date clearly yields better results for both performance measures. For the 1-M test cases, for instance, the impact on revenue varies between 0.27% and 1.74%. In fact, revenue gains in this range are of major economic importance, as an increase in revenue is often transformed almost one-to-one to profit because no or only negligible costs are necessary to realize it. Furthermore, as additionally illustrated in Fig. 4 in respect of 1-M and BP-DLP, the impact of the notification date is obviously stronger for less reliable forecasts, underlining the intuition that late reallocation is especially helpful in situations of rather uncertain demand. Finally, in a comparison of the three controls, CEC always experiences the biggest revenue gains.




Fig. 4 Impact of the notification date \overline{OG}^s (1-M, BP-DLP)

Table 5 \overline{OG}^s and \overline{CU}^s for different notification dates (days before departure)

	1-M						2-M					
	CEC		BP-DLP		BP-RLP		CEC		BP-DLP		BP-RLP	
	\overline{OG}	\overline{CU}	\overline{OG}	\overline{CU}	\overline{OG}	\overline{CU}	\overline{OG}	\overline{CU}	\overline{OG}	\overline{CU}	\overline{OG}	\overline{CU}
$\delta_i = 0$												
40 days	-2.18	98.48	-3.14	98.97	-2.48	98.18	-1.60	96.01	-2.47	96.25	-1.87	95.63
20 days	-1.92	98.72	-2.80	99.21	-2.21	98.38	-1.53	96.09	-2.34	96.34	-1.80	95.69
0 days	-1.42	99.04	-2.62	99.30	-2.21	98.38	-1.34	96.23	-2.25	96.40	-1.80	95.67
$\delta_i = 0.25$												
40 days	-2.85	97.94	-3.74	98.46	-3.47	97.37	-2.33	95.58	-3.03	95.79	-2.82	94.90
20 days	-2.51	98.22	-3.34	98.72	-3.15	97.58	-2.18	95.70	-2.85	95.92	-2.68	95.00
0 days	-2.07	98.50	-3.14	98.83	-3.14	97.59	-1.96	95.85	-2.73	95.99	-2.64	95.00
$\delta_i = 0.5$												
40 days	-4.76	96.38	-5.43	96.91	-5.51	95.56	-3.68	94.31	-4.31	94.49	-4.51	93.26
20 days	-3.74	97.16	-4.23	97.72	-4.65	96.15	-3.37	94.52	-3.94	94.71	-4.25	93.42
0 days	-3.02	97.60	-3.92	97.89	-4.55	96.19	-3.16	94.66	-3.82	94.78	-4.19	93.43

5 Summary and conclusion

In this paper, we have introduced new model formulations for revenue management with flexible products. They allow for explicitly incorporating arbitrary notification dates. We extended several popular control mechanisms, like booking limits and bid-price controls, enabling the firm to take advantage of the flexibility between the sale and notification date. An extensive simulation study was conducted to analyze the performance of the controls in respect of various test cases from passenger airline revenue management. Imprecise and biased demand forecasts were especially emphasized. In this context, we were able to isolate the benefits resulting from the flexibility. Our results confirm the intuition that flexibility is more important to the firm if uncertainty is higher and that a late notification date helps to further increase revenues. In scenarios based on realistic assumptions, we observed an increase in revenue of up to 4%, leading to a tremendous gain in profit in fields of application that typically use revenue management.

Overall, we conclude that the potential of flexible products goes beyond marketing effects like demand induction. When introducing such products, firms should not limit themselves to ad-hoc variants of capacity control which imply that the execution mode is irrevocably determined immediately after sale. Instead, they should carefully consider implementing adequate controls to benefit from a later notification date, following the more versatile approaches introduced in this paper.

Appendix

A Proof of acceptance criteria (3.18)–(3.20)

In the following, we show that the allocation y_{jm}^{e*} in respect of flexible product $j \in \mathcal{J}$ in mode $m \in \mathcal{M}_j$ in an optimal solution of model (3.6)–(3.12) is positive if and only if conditions (3.18)–(3.20) hold.

The dual of model (3.6)–(3.12) is given by

$$\tilde{V}D^{DLP}(t, \mathbf{c}, \mathbf{Y}^a) = \min_{\mathbf{v}, \mathbf{w}, \mathbf{u}, \mathbf{o}} \left(\sum_{h \in \mathcal{H}} c_h \cdot v_h + \sum_{i \in \mathcal{I}} \bar{D}_{it}^s \cdot w_i + \sum_{j \in \mathcal{J}} \bar{D}_{jt}^f \cdot u_j + \sum_{j \in \mathcal{J}} Y_j^a \cdot o_j \right) \quad (\text{A.1})$$

subject to

$$\sum_{h \in \mathcal{H}} a_{hi} \cdot v_h + w_i \geq r_i \quad \text{for all } i \in \mathcal{I}, \quad (\text{A.2})$$

$$\sum_{h \in \mathcal{H}} a_{hm} \cdot v_h + u_j \geq f_{jm} \quad \text{for all } j \in \mathcal{J}, m \in \mathcal{M}_j, \quad (\text{A.3})$$

$$\sum_{h \in \mathcal{H}} a_{hm} \cdot v_h + o_j \geq f_{jm} \quad \text{for all } j \in \mathcal{J}, m \in \mathcal{M}_j, \quad (\text{A.4})$$

$$v_h \geq 0 \quad \text{for all } h \in \mathcal{H}, \quad (\text{A.5})$$

$$w_i \geq 0 \quad \text{for all } i \in \mathcal{I}, \quad (\text{A.6})$$

$$u_j \geq 0 \quad \text{for all } j \in \mathcal{J}, \quad (\text{A.7})$$

$$o_j \in \mathbb{R} \quad \text{for all } j \in \mathcal{J}, \quad (\text{A.8})$$

with $\mathbf{v} = (v_1, \dots, v_l)$, $\mathbf{w} = (w_1, \dots, w_{n^s})$, $\mathbf{u} = (u_1, \dots, u_{n^f})$, and $\mathbf{o} = (o_1, \dots, o_{n^f})$.

Let $(\mathbf{x}^*, \mathbf{y}^{e*}, \mathbf{y}^{a*})$ and $(\mathbf{v}^*, \mathbf{w}^*, \mathbf{u}^*, \mathbf{o}^*)$ be optimal solutions of the primal and the dual problem, respectively. Then the following complementary slackness conditions hold (among others):

$$y_{jm}^{e*} \cdot \left(-f_{jm} + \sum_{h \in \mathcal{H}} a_{hm} \cdot v_h^* + u_j^* \right) = 0 \quad \text{for all } j \in \mathcal{J}, m \in \mathcal{M}_j. \quad (\text{A.9})$$

Let us now consider constraint (A.3), which can be rewritten as $u_j \geq f_{jm} - \sum_{h \in \mathcal{H}} a_{hm} \cdot v_h$ for all $j \in \mathcal{J}$, $m \in \mathcal{M}_j$. As the dual is a minimization problem and $\sum_{j \in \mathcal{J}} \bar{D}_{jt}^f \cdot u_j$ with $\bar{D}_{jt}^f > 0$ for all $j \in \mathcal{J}$ is part of the objective function, it follows that

$$u_j^* = \max \left\{ \max_{m \in \mathcal{M}_j} \left\{ f_{jm} - \sum_{h \in \mathcal{H}} a_{hm} \cdot v_h^* \right\}, 0 \right\} \quad \text{for all } j \in \mathcal{J}. \quad (\text{A.10})$$

We now distinguish the following three (exhaustive) cases in respect of a certain flexible product $j \in \mathcal{J}$:

Case 1:
$$\max_{m \in \mathcal{M}_j} \left\{ f_{jm} - \sum_{h \in \mathcal{H}} a_{hm} \cdot v_h^* \right\} < 0.$$

From (A.10), it follows that $u_j^* = 0$ and from constraint (A.3) that $f_{jm} < \sum_{h \in \mathcal{H}} a_{hm} \cdot v_h^*$

for all $m \in \mathcal{M}_j$. From complementary slackness (A.9), we can conclude that $y_{jm}^{e*} = 0$ for all $m \in \mathcal{M}_j$, meaning that there is no positive contingent for flexible product j in any of its possible execution modes.

Case 2:
$$\max_{m \in \mathcal{M}_j} \left\{ f_{jm} - \sum_{h \in \mathcal{H}} a_{hm} \cdot v_h^* \right\} = 0.$$

Again, it follows from (A.10) that $u_j^* = 0$. We now partition the set \mathcal{M}_j into two subsets $\mathcal{M}_j^{(1)}$ and $\mathcal{M}_j^{(2)}$:

2a) There is a non-empty set $\mathcal{M}_j^{(1)}$ of modes without slackness in constraint (A.3). We can conclude that $f_{jm} = \sum_{h \in \mathcal{H}} a_{hm} \cdot v_h^*$ and so $y_{jm}^{e*} > 0$ by (A.9) for all $m \in \mathcal{M}_j^{(1)}$, if the optimal solution is not degenerated.

2b) The set $\mathcal{M}_j^{(2)}$ contains modes with slackness in constraint (A.3) so that $f_{jm} < \sum_{h \in \mathcal{H}} a_{hm} \cdot v_h^*$ and so $y_{jm}^{e*} = 0$ by (A.9) for all $m \in \mathcal{M}_j^{(2)}$.

Case 3:
$$\max_{m \in \mathcal{M}_j} \left\{ f_{jm} - \sum_{h \in \mathcal{H}} a_{hm} \cdot v_h^* \right\} > 0.$$

From (A.10), it follows that $u_j^* > 0$ and so $u_j^* = \max_{m \in \mathcal{M}_j} \left\{ f_{jm} - \sum_{h \in \mathcal{H}} a_{hm} \cdot v_h^* \right\}$. Again, we

distinguish two sets $\mathcal{M}_j^{(1)}$ and $\mathcal{M}_j^{(2)}$ representing a partition of \mathcal{M}_j :

- 3a) The non-empty set $\mathcal{M}_j^{(1)}$ contains the modes m with $f_{jm} - \sum_{h \in \mathcal{H}} a_{hm} \cdot v_h^* = \max_{m' \in \mathcal{M}_j} \left\{ f_{jm'} - \sum_{h \in \mathcal{H}} a_{hm'} \cdot v_h^* \right\}$. By (A.9), we can conclude that $y_{jm}^{e*} > 0$ for all $m \in \mathcal{M}_j^{(1)}$ (as $u_j^* = f_{jm} - \sum_{h \in \mathcal{H}} a_{hm} \cdot v_h^*$), if the optimal solution is not degenerated.
- 3b) For all $m \in \mathcal{M}_j^{(2)}$, $f_{jm} - \sum_{h \in \mathcal{H}} a_{hm} \cdot v_h^* < \max_{m' \in \mathcal{M}_j} \left\{ f_{jm'} - \sum_{h \in \mathcal{H}} a_{hm'} \cdot v_h^* \right\}$ so that $y_{jm}^{e*} = 0$ by (A.9) (as $u_j^* > f_{jm} - \sum_{h \in \mathcal{H}} a_{hm} \cdot v_h^*$).

To complete the proof, one can now easily verify that the optimal allocation y_{jm}^{e*} in respect of flexible product $j \in \mathcal{J}$ in mode $m \in \mathcal{M}_j$ is positive if conditions (3.18)–(3.20) hold (cases 2a and 3a with the bid-price π_h set to v_h^*). Otherwise, the optimal allocation will be zero (cases 1, 2b and 3b). Note that the result is only valid for non-degenerated optimal solutions. In case of degeneration, the control suffers the same defects as they are reported for standard revenue management bid-price controls (see, e.g., Bertsimas and Popescu 2003).

B Network 1

Table B-1 LF_h in Network 1

Nominal load factor LF_h for leg h	1	2	3	4	Mean
1-H - high demand	1.5	1.7	1.3	1.1	1.4
1-M - medium demand	1.2	1.4	0.8	0.9	1.1
1-L - low demand	1.1	0.7	0.9	1.05	0.9
1-E - erratic demand	0.6	1.5	0.9	1.1	1.0

C Network 2

The calculation of the expected demands \bar{D}_{iT}^s and \bar{D}_{jT}^f for all specific products $i \in \mathcal{I}$ as well as for flexible products $j \in \mathcal{J}$, based on the demand situation, is outlined in the following. To simplify the notation, we partition the set of resources \mathcal{H} and create three subsets $\mathcal{H}_S = \{1, \dots, 6\}$, $\mathcal{H}_M = \{7, \dots, 13\}$, and $\mathcal{H}_L = \{14, 15\}$, which contain the short, medium, and long haul flights, respectively.

As demand \bar{D}_{jT}^f for a flexible product equals $\rho = 20\%$ of the total demand for itineraries on which it is based, demand \bar{D}_{iT}^f for the product “direct flight from A to D with flexible departure time” is given by $\bar{D}_{iT}^f = \sum_{l \in \mathcal{H}_L} \rho \cdot LF_l \cdot C_l$ (see Table C-1

in respect of LF_l and C_l). The remaining demand on long haul flight $l \in \mathcal{H}_L$ is for specific products: $d_l^s = (1 - \rho) \cdot LF_l \cdot C_l$.

For short haul flights $s \in \mathcal{H}_S$ and medium haul flights $m \in \mathcal{H}_M$ the definition is complicated by the necessity to consider connecting itineraries. Demand d_{sm} for an itinerary, consisting of a short haul flight $s \in \mathcal{H}_S$ and a medium haul flight $m \in \mathcal{H}_M$, is a proportion of total demand $LF_m \cdot C_m$ for the medium haul leg m : $d_{sm} = \xi_{sm} \cdot \varphi_m \cdot LF_m \cdot C_m$ for all $s \in \mathcal{H}_S$ and $m \in \mathcal{H}_M$. The value φ_m (see Table C-1) denotes a proportion of the demand for connecting itineraries on the medium haul flight m . Since some medium haul flights have the possibility to connect to or from various short haul flights s , ξ_{sm} defines the proportion of connecting passengers on flight m also flying on s (Table C-2 omits the ξ_{sm} set to zero corresponding to itineraries not offered). Thus, the demand for itineraries consisting only of a short or medium haul flight leg – where only specific products are offered – is given by $d_s^s = LF_s \cdot C_s - \sum_{m \in \mathcal{H}_M} d_{sm}$ and $d_m^s = (1 - \varphi_m) \cdot LF_m \cdot C_m = LF_m \cdot C_m - \sum_{s \in \mathcal{H}_S} d_{sm}$, respectively.

Demand \bar{D}_{2T}^f for the second flexible product “one-stop flight from A to D, early departure” is again a proportion of the demand on the corresponding itineraries and given by $\bar{D}_{2T}^f = \rho \cdot (d_{4,7} + d_{5,7} + d_{1,10} + d_{1,11})$. Similarly, $\bar{D}_{3T}^f = \rho \cdot (d_{6,9} + d_{3,12} + d_{3,13})$ is the demand for the third, “one-stop flight from A to D, late departure.” The demand for the specific products of a connecting itinerary d_{sm}^s equals $d_{sm}^s = (1 - \rho) \cdot d_{sm}$ if the itinerary is also part of the second or third flexible product and $d_{sm}^s = d_{sm}$ otherwise. Finally, the itinerary demand for specific products (d_s^s , d_m^s , d_l^s , or d_{sm}^s , respectively) is assigned to the four booking classes according to the ratios defined in Section 4.1.

Table C-1 C_h , φ_m and LF_h in Network 2

Leg	h	C_h	φ_m	Nominal load factor LF_h for leg h			
				2-H	2-M	2-L	2-E
AB	9:00	1	250	1.4	0.8	0.7	1.4
	10:00	2	250	1.3	1.3	1.1	0.7
	13:00	3	250	1.2	1.1	1.0	1.2
CD	9:30	4	200	1.3	0.9	0.8	0.8
	10:30	5	250	1.5	1.2	1.1	1.5
	21:00	6	200	1.1	1.1	0.9	1.1
AC	7:00	7	400	0.4	1.2	0.9	0.7
	8:00	8	400	0.4	1.4	1.2	1.0
	18:30	9	300	0.4	1.2	1.1	0.9
BD	12:00	10	300	0.4	1.1	1.0	1.2
	13:00	11	400	0.4	1.6	1.4	0.8
	16:00	12	400	0.3	1.0	1.1	1.1
	17:00	13	300	0.3	1.3	0.9	0.7
AD	6:00	14	400		1.2	0.85	0.85
	17:00	15	400		1.4	1.3	1.1

Table C-2 ξ_{sm} (proportion of connecting passengers on flight m also flying on s)

	A-C			B-D			
	7:00	8:00	18:30	12:00	13:00	16:00	17:00
C-D 9:30	0.8			A-B 9:00	1	0.2	
C-D 10:30	0.2			A-B 10:00		0.8	
C-D 21:00		1	1	A-B 13:00			1 1

References

- Alstrup J, Boas S, Madsen OBG, Vidal RVV (1986) Booking policy for flights with two types of passengers. *Eur J Oper Res* 27:274–288
- Bartodziej P, Derigs U (2004) On an experimental algorithm for revenue management for cargo airlines. In: Ribeiro CC, Martins SL (eds) *Experimental and efficient algorithms*. Springer, Berlin, pp 57–71
- Bartodziej P, Derigs U, Zils M (2007) O&D revenue management in cargo airlines—A mathematical programming approach. *OR Spectrum* 29:105–121
- Belobaba PP (1987) Airline yield management—An overview of seat inventory control. *Transport Sci* 21:63–73
- Bertsekas DP (2005) *Dynamic programming and optimal control*, vol 1, 3rd edn. Athena, Belmont
- Bertsimas D, Popescu I (2003) Revenue management in a dynamic network environment. *Transport Sci* 37:257–277
- Büke B, Yildirim U, Kuyumcu HA (2008) New stochastic linear programming approximations for network capacity control problem with buy-ups. *J Revenue Pricing Manage* 7:61–84
- Chen VCP, Günther D, Johnson EL (2003) Routing considerations in airline yield management. In: Ciriani TA, Fasano G, Gliozzi S, Tadei R (eds) *Operations research in space and air*. Kluwer, Dordrecht, pp 333–350

- Chiang W-C, Chen JCH, Xu X (2007) An overview of research on revenue management—Current issues and future research. *Int J Revenue Manage* 1:97–128
- De Boer SV, Freling R, Piersma N (2002) Mathematical programming for network revenue management revisited. *Eur J Oper Res* 137:72–92
- Gallego G, Iyengar G, Phillips R, Dubey A (2004) Managing flexible products on a network. Working Paper TR-2004-01, Columbia University, New York
- Gallego G, Phillips R (2004) Revenue management of flexible products. *Manuf Serv Oper Manage* 6:321–337
- Gallien J, Le Tallec Y, Schoenmeyr T (2004) A model for make-to-order revenue management. Working Paper, MIT, Cambridge
- Geraghty MK, Johnson E (1997) Revenue management saves national car rental. *Interfaces* 27(1):107–127
- Jalora A (2006) Order acceptance and scheduling at a make-to-order system using revenue management. Ph.D. Thesis, Texas A&M University, College Station
- Karaesmen IZ (2001) Three essays on revenue management. Ph.D. Thesis, Columbia University, New York
- Karaesmen IZ, van Ryzin G (2004) Overbooking with substitutable inventory classes. *Oper Res* 52:83–104
- Kimes SE (2000) A strategic approach to yield management. In: Ingold A, McMahon-Beattie U, Yeoman I (eds) *Yield management—Strategies for the service industries*, 2nd edn. Thomson, London, pp 3–14
- Kimms A, Müller-Bungart M (2007a) Revenue management for broadcasting commercials—The channel's problem of selecting and scheduling advertisements to be aired. *Int J Revenue Manage* 1:28–44
- Kimms A, Müller-Bungart M (2007b) Simulation of stochastic demand data streams for network revenue management problems. *OR Spectrum* 29:5–20
- Klein R (2007) Network capacity control using self-adjusting bid-prices. *OR Spectrum* 29:39–60
- Law AM (2007) *Simulation modeling and analysis*, 4th edn. McGraw-Hill, Boston
- Lee AO (1990) Airline reservations forecasting—Probabilistic and statistical models of the booking process. Ph.D. Thesis, MIT, Cambridge
- McGill JI, van Ryzin GJ (1999) Revenue management—Research overview and prospects. *Transport Sci* 33:233–256
- Netessine S, Dobson G, Shumsky RS (2002) Flexible service capacity—Optimal investment and the impact of demand correlation. *Oper Res* 50:375–388
- Phillips RL (2005) *Pricing and revenue optimization*. Stanford University Press, Stanford
- Shumsky RA, Zhang F (2009) Dynamic capacity management with substitution. *Oper Res* (forthcoming). doi:10.1287/opre.1080.0610
- Simpson RW (1989) Using network flow techniques to find shadow prices for market and seat inventory control. MIT flight transportation laboratory memorandum M89-1, Cambridge
- Smith BC, Penn CW (1988) Analysis of alternative origin-destination control strategies. Proceedings of the twenty eighth annual AGIFORS symposium, New Seabury

- Spengler T, Rehkopf S, Volling T (2007) Revenue management in make-to-order manufacturing—An application to the iron and steel industry. *OR Spectrum* 29:157-171
- Stadtler H (2008) Supply chain management—An overview. In: Stadtler H, Kilger C (eds) *Supply chain management and advanced planning—Concepts, models, software, and case studies*. Springer, Berlin, pp 1–36
- Talluri KT, van Ryzin GJ (1998) An analysis of bid-price controls for network revenue management. *Manage Sci* 44:1577–1593
- Talluri KT, van Ryzin GJ (1999) A randomized linear programming method for computing network bid prices. *Transport Sci* 33:207–216
- Talluri KT (2001) Airline revenue management with passenger routing control—A new model with solution approaches. *Int J Serv Tech Manage* 2:102–115
- Talluri KT, van Ryzin GJ (2004) *The theory and practice of revenue management*. Kluwer, Boston
- Topaloglu H (2009) On the asymptotic optimality of the randomized linear program for network revenue management. *Eur J Oper Res* 197:884–896
- Tscheulin DK, Lindenmeier J (2003) Yield-Management—Ein State-of-the-Art. *Z Betriebswirt* 73:629–662
- Weatherford LR, Bodily SE (1992) A taxonomy and research overview of perishable-asset revenue management—Yield management, overbooking, and pricing. *Oper Res* 40:831–844
- Williamson EL (1992) *Airline network seat control*. Ph.D. Thesis, MIT, Cambridge