Practical Decision Rules for Risk-Averse Revenue Management using Simulation-Based Optimization

February 05, 2016

Sebastian Koch

Chair of Analytics & Optimization, University of Augsburg,
Universitätsstraße 16, 86159 Augsburg, Germany

e-mail: sebastian.koch@wiwi.uni-augsburg.de

Jochen Gönsch (✉)

Chair of Service Operations, Mercator School of Management, University of Duisburg-Essen,
Lotharstraße 65, 47057 Duisburg, Germany

Phone: +49 203 379 4369
Fax: +49 203 379 1760
e-mail: jochen.goensch@uni-due.de

Michael Hassler

Chair of Analytics & Optimization, University of Augsburg,
Universitätsstraße 16, 86159 Augsburg, Germany

e-mail: michael.hassler@wiwi.uni-augsburg.de

Robert Klein

Chair of Analytics & Optimization, University of Augsburg,
Universitätsstraße 16, 86159 Augsburg, Germany

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

Sebastian Koch is a research and teaching assistant at the Faculty of Business and Economics at the University of Augsburg. His research focuses on revenue management and approximate dynamic programming and has been published in journals such as Computers & Operations Research and International Journal of Production Economics.

Jochen Gönsch is professor for Service Operations at the Mercator School of Management at the University of Duisburg-Essen. He obtained his doctorate from the University of Augsburg for a thesis on revenue management with flexibility in the underlying service/production systems. Currently, he researches operations management and pricing with a focus on mobility, remanufacturing, and renewable energy.

Michael Hassler is a research and teaching assistant at the Faculty of Business and Economics at the University of Augsburg. His research focuses on approximate dynamic programming with the application areas of revenue management and renewable energy.

Robert Klein holds the Chair of Analytics & Optimization at the Faculty of Business and Economics at the University of Augsburg. He obtained his doctorate from the Darmstadt University of Technology. Robert Klein's research addresses practical decision support using quantitative methods with a focus on pricing and revenue management. He has worked, for example, with IBM, Lufthansa Systems, and Sixt.

Practical Decision Rules for Risk-Averse Revenue Management by means of Simulation-Based Optimization

February 05, 2016

Abstract

In practice, human-decision makers often feel uncomfortable with the risk-neutral revenue management systems' output. Reasons include a low number of repetitions of similar events, a critical impact of the achieved revenue for economic survival, or simply business constraints imposed by management. However, solving capacity control problems is a challenging task for many risk measures and the approaches are often not compatible with existing software systems.

In this paper, we propose a flexible framework for risk-averse capacity control under customer choice behavior. Existing risk-neutral decision rules are augmented by the integration of adjustable parameters. Our key idea is the application of simulation-based optimization (SBO) to calibrate these parameters. This allows to easily tailor the resulting capacity control mechanism to almost every risk measure and customer choice behavior.

In an extensive simulation study, we analyze the impact of our approach on expected utility, conditional value-at-risk (CVaR), and expected value. The results show a superior performance in comparison to risk-neutral approaches from literature.

Keywords: Revenue Management, Capacity Control, Risk-Aversion, Conditional Value-at-Risk

INTRODUCTION

During the last decades, revenue management has become one of the most successful fields of application for operations research in practice. Its main task is capacity control, which is usually described as controlling the availability of differentiated products over a given booking horizon such that the expected revenue is maximized. The assumption of risk-neutrality lies at the heart of this classical definition and is justified by a large number of repetitions of similar decision problems. However, human decision makers, who tend to be risk-averse, often doubt this assumption. In daily practice, they feel uncomfortable with the capacity control system's output and overwrite it manually with less aggressive decisions. Furthermore, in many fields of application, the number of repetitions is too small to justify the use of expected value and a single event is critical for economic survival. Risk-aversion first became popular in economics and finance, but it is today also increasingly considered in revenue management. The underlying trade-off is to give up a portion of expected value in order to reduce the risk of poor outcomes.

The problem of risk-averse capacity control can be solved to optimality by dynamic programming (DP). However, building a DP formulation is a challenging task for many risk measures. In many cases, the state space must be augmented and the resulting DP formulation becomes intractable. Furthermore, DP formulations are not compatible with many existing revenue management systems.

Our main contribution is to propose a flexible framework for risk-averse capacity control. In practice, revenue management systems are fixed in the long run and the capacity control process is modeled by standard decision rules such as bid prices. Therefore, our framework is based on the risk-neutral formulation. Risk-aversion is then integrated by augmenting existing capacity control mechanisms with a few parameters that can be calibrated. Existing research recommends that this is done manually by human decision makers. However, we suggest the use of simulation-based optimization (SBO) which allows automated optimization and a higher number of parameters. The resulting approach is quite general. It can be used with arbitrary demand models, risk measures, and network structures. In an extensive simulation study, we illustrate the impact of our approach on expected utility, conditional value-at-risk (CVaR), and expected revenue in various settings with customer choice and different network structures.

The remainder of this paper is structured as follows: First, we restate the risk-neutral problem of capacity control under customer choice behavior, review the relevant scientific literature and position our work. Based upon this, we present our framework for risk-averse capacity control, including a detailed description of the components. We continue with the simulation study, followed by a discussion of the results and a conclusion.

BACKGROUND AND PREVIOUS RESEARCH

Research from three areas of revenue management is relevant for our work. First, we restate the problem of risk-neutral capacity control under customer choice behavior and summarize standard solution approaches. Then, we discuss research on risk-averse capacity control and the use of SBO for capacity control.

Risk-neutral capacity control under customer choice behavior

Initially, revenue management (RM) was based on the well-known independent demand assumption. Overviews can be found in the textbooks of Talluri and van Ryzin (2004b) and Phillips (2005).

Later, research considered that most customers actually choose between several more or less suitable products. Gallego et al. (2004), Talluri and van Ryzin (2004a), and Liu and van Ryzin (2008) established capacity control under a general discrete choice model of demand. In this setting, a firm disposes of resources $i = 1, \ldots, m$ which are jointly used by products $j = 1, \ldots, n$. The products are associated with revenues $\boldsymbol{r} = (r_1, \ldots, r_n)^T$. Furthermore, each product j has a capacity consumption $\boldsymbol{a}_j = \left(a_{1j}, \ldots, a_{mj}\right)^T$, which is either $a_{ij} = 1$ if product j requires resource i or $a_{ij} = 0$ else. Resources' remaining capacity is denoted by the vector $\boldsymbol{c} = (c_1, \ldots, c_m)^T$, the initial endowment is given by $\boldsymbol{c}^0 = (c_1^0, \ldots, c_m^0)^T$. Customers arrive successively and stochastically over time. The booking horizon is discretized into sufficiently small time periods $t = 1, \ldots, T$, such that in each period t at most one customer arrives. Thus, at most one product can be sold in each period. The periods are numbered forward in time. Any capacity remaining at the end of the booking horizon is worthless and overbooking of the given resources' capacity is not allowed.

In each period t, the firm's risk-neutral decision problem is to determine a subset of products to offer, called the offer set, so that the overall expected revenue $V_1(\boldsymbol{c}^0)$ is maximized. The offer set is captured by the vector $\boldsymbol{x} = (x_1, \ldots, x_n)^T$ of binary decision variables with $x_j = 1$ if product j is offered for sale. Product j is sold with probability $p_{tj}(\boldsymbol{x})$ and no purchase is made with probability $p_{t0}(\boldsymbol{x})$.

Let the value function $V_t(\boldsymbol{c})$ denote the optimal expected revenue-to-go in period t with capacity $\boldsymbol{c}$ and let $\Delta_j V_t(\boldsymbol{c}) \coloneqq V_t(\boldsymbol{c}) - V_t(\boldsymbol{c} - \boldsymbol{a}_j)$ denote the opportunity cost of selling one unit of product j. Then, $V_t(\boldsymbol{c})$ and the expected revenue-maximizing offer set can be computed recursively by the following DP formulation (DP-EV)

$$V_t(\boldsymbol{c}) = \max_{\boldsymbol{x}} \left\{ \sum_{j=1}^{n} p_{tj}(\boldsymbol{x}) \cdot \left(r_j - \Delta_j V_{t+1}(\boldsymbol{c}) \right) \right\} + V_{t+1}(\boldsymbol{c}) \tag{1}$$

subject to the boundary conditions $V_t(\boldsymbol{c}) = -\infty$ if $\boldsymbol{c} \ngeq \boldsymbol{0}$ and $V_{T+1}(\boldsymbol{c}) = 0$ if $\boldsymbol{c} \geq \boldsymbol{0}$.

Two issues render DP-EV difficult to solve optimally: recursively calculating the opportunity cost $\Delta_j V_t(\boldsymbol{c})$ and solving the maximization over all 2^n possible offer sets. Over time, different heuristic approaches have been developed. Regarding the first issue, virtually all approaches use additive bid prices π_{tic_i} that reflect the current value of one unit of capacity of resource i in period t with remaining capacity c_i. With these values, an approximation $\widetilde{\Delta}_j V_t(\boldsymbol{c})$ of the opportunity cost can be obtained:

$$\widetilde{\Delta}_j V_t(\boldsymbol{c}) = \sum_{i=1}^{m} a_{ij} \cdot \pi_{tic_i} \tag{2}$$

The approaches differ in how the bid prices are computed, but the main idea is usually to derive an easy-to-compute upper bound on $V_t(\boldsymbol{c})$ and use information from this upper bound to approximate the opportunity cost in an offline stage (that is, before the booking horizon starts). Such approximations can be found, for example, in Liu and van Ryzin (2008), Miranda Bront et al. (2009), Zhang and Adelman (2009), and Meissner and Strauss (2012b). Online (that is, during the booking horizon), the offer set is then determined by solving the maximization, which is an assortment optimization problem (see, e.g., Miranda Bront et al. (2009)):

$$\max_{\boldsymbol{x}} \left\{ \sum_{j=1}^{n} p_{tj}(\boldsymbol{x}) \cdot \left(r_j - \widetilde{\Delta}_j V_{t+1}(\boldsymbol{c}) \right) \right\} \tag{3}$$

The technique used to solve (3) strongly depends on the choice model assumed. For example, under the independent demand model, (3) reduces to the classical method of simply offering all products for which revenue exceeds (an approximation of) opportunity cost:

$$r_j \geq \widetilde{\Delta}_j V_{t+1}(\boldsymbol{c}) \tag{4}$$

A popular way to manage the selling process, in particular in practice, is to use this independent demand decision rule (4) in combination with additive bid prices (2). This kind of capacity control approach is often referred to as bid price control. Even if demand is not independent, decision rule (4) can be used heuristically (see, e.g., Chaneton and Vulcano (2011) and Meissner and Strauss (2012a)).

Table 1 summarizes the notation used throughout this section.

$i = 1, \dots, m$	resources	$p_{tj}(\boldsymbol{x})$	purchase probability of product j given offer set $\boldsymbol{x}$
$j = 1, \dots, n$	products	$p_{t0}(\boldsymbol{x})$	no-purchase probability given offer set $\boldsymbol{x}$
$t = 1, \dots, T$	time periods (numbered forward)	$V_t(\boldsymbol{c})$	optimal expected revenue-to-go in period t with capacity $\boldsymbol{c}$
$\boldsymbol{c} = (c_1, \dots, c_m)^T$	remaining capacity	$\Delta_j V_t(\boldsymbol{c})$	opportunity cost of product j
$\boldsymbol{c}^0 = (c_1^0, \dots, c_m^0)^T$	initial capacity	$\widetilde{\Delta}_j V_t(\boldsymbol{c})$	approximation of opportunity cost of product j
$\boldsymbol{r} = (r_1, \dots, r_n)^T$	product revenues	π_{tic_i}	bid price of resource i in period t with remaining capacity c_i
$\boldsymbol{a}_j = \left(a_{1j}, \dots, a_{mj}\right)^T$	capacity consumption of product j		
$\boldsymbol{x} = (x_1, \dots, x_n)^T$	offer set of products		

Table 1: Notation introduced in this section

Risk-averse capacity control

In this section, we briefly outline the consideration of risk in the academic literature on capacity control. Only the most relevant literature is mentioned. For a recent review, we refer to Gönsch and Hassler (2014) and the references therein.

The need for considering risk-aversion in capacity control was first raised by Lancaster (2003) who proposed a risk-adjusted revenue per available seat mile. Weatherford (2004) then modified the famous EMSR-b heuristic of Belobaba (1992) by substituting revenues with a risk-averse utility function. Barz (2007), Barz and Waldmann (2007), and Feng and Xiao (2008) use an exponential utility function to model risk-aversion, but instead of altering a heuristic, they work with the original DP formulation. Assuming independent demand and a single resource, they show that several well-known properties regarding the structure of an optimal policy carry over from the risk-neutral case. In addition, Barz (2007) extends this analysis to the case of customer choice behavior. Zhuang and Li (2011) examine optimal booking limits with an atemporal utility function to address risk-aversion. Furthermore, there are two publications from Koenig and Meissner (Koenig and Meissner (2015b, 2015c)) who consider target percentile risk and value-at-risk.

Most relevant to our work are Huang and Chang (2011) and Koenig and Meissner (2015a). Similar to our work, they modify existing capacity control approaches to address risk-aversion. In particular, Huang and Chang (2011) heuristically consider risk-aversion via a discount factor on the opportunity cost in the DP formulation. This factor is either constant or a function of remaining demand and capacity. Koenig and Meissner (2015a) extend this analysis. In addition, they consider a discount factor on the opportunity cost from a risk-neutral DP formulation and an alternative function of demand and capacity. However, the approaches are restricted to a few parameters that are calibrated manually. Moreover, as in all the literature on risk-averse capacity control so far, only single-leg settings are considered.

Finally, the literature on risk-averse dynamic pricing is related to us (see Gönsch et al. (2015) for a recent review). The main difference between dynamic pricing and capacity control is that the decision maker influences demand by setting the prices of products instead of choosing the offer set, while the general setting is quite similar (see, e.g., Gallego and van Ryzin (1997) and Talluri and van Ryzin (2004b) for problems with risk-neutral decision makers). Thus, the incorporation of risk-aversion is done in a similar fashion (see, e.g., Li and Zhuang (2009) for utility functions; Feng and Xiao (1999) for revenue variance; Levin et al. (2008) for target percentile risk; Gönsch et al. (2015) for conditional value-at-risk).

Simulation-based optimization for capacity control

Until now, SBO has only been used in risk-neutral capacity control. For a general overview of SBO please refer to, for example, Gosavi (2015) or Spall (2003). Robinson (1995) was the first to use SBO in the context of revenue management to approximate the optimal booking limit policy in the single-leg case. More recent research derives stochastic gradients of the value function and uses estimates of these gradients in the optimization step. Bertsimas and de Boer (2005) present an algorithm for the improvement of booking limits, which uses a discretization of the state space for value function estimation. Gosavi et al. (2007) show that an algorithm based on simultaneous perturbation for the improvement of booking limits outperforms both EMSR-b and DAVN-EMSR-b for single-leg and network problems, respectively. Topaloglu (2008) and van Ryzin and Vulcano (2008b) improve bid prices and nested protection limits, respectively, by using a continuous approximation of the discrete problem which enables an exact recursive computation of gradients. While all the aforementioned papers follow the independent demand assumption, van Ryzin and Vulcano (2008a) use a procedure similar to van Ryzin and Vulcano (2008b) in order to improve nested protection limits under customer choice behavior. Chaneton and Vulcano (2011) present a stochastic gradient algorithm for improvement of bid prices with customer choice.

However, the use of stochastic gradients is feasible only if a recursive formulation of the value function is available. Unfortunately, this is often not the case or very challenging for the objective functions considered in risk-averse revenue management. Therefore, we concentrate on purely numerical approaches in our paper that can be adapted easily to different objectives. In this sense, Klein (2007) is closest to us. He introduces auto-adaptive bid prices by means of the metaheuristic scatter search assuming independent demand.

RISK-AVERSE CAPACITY CONTROL USING SBO

In this section, we first present an overview of the new framework allowing the incorporation of risk-aversion. Then, we turn to the most important components and describe in detail the risk measures and capacity control approaches considered in this study. Table 2 summarizes the additional notation introduced in the following section.

n^{calib}	number of calibration streams	$U_\gamma(\cdot)$	exponential utility function with level of risk-aversion γ
n^{eval}	number of evaluation streams	$CVaR_\alpha(\cdot)$	Conditional Value-at-Risk (CVaR) at probability level α
R	revenue obtained	$\boldsymbol{\theta}$	(arbitrary) parameters to integrate risk-aversion into capacity control mechanisms
$F(y)$	distribution of total revenue R, i.e $F(y) = \mathbb{P}(R \leq y)$		
$U(\cdot)$	utility function		

Table 2: Notation introduced in this section

Overview

Our basic idea is to modify standard approaches appropriately to account for risk-aversion. This modification is governed by parameters $\boldsymbol{\theta}$, which are determined by an out-of-the-box iterative SBO algorithm before the beginning of the booking horizon. The whole process consists of three steps (see Figure 1).

The optimization step aims at improving values for the parameters $\boldsymbol{\theta}$. It passes tentative values to the simulation step to estimate their performance. The simulation step in turn mimics sales processes using n^{calib} independent customer demand streams, each encompassing the whole booking horizon. This calibration set is generated in advance according to the firm's belief about future demand. For each demand stream, the control mechanism with the current values of $\boldsymbol{\theta}$ is applied and a total per-stream revenue is obtained. All per-stream revenues are used to calculate a risk measure, which is passed back to the optimization step as an estimate of the $\boldsymbol{\theta}$-values' performance. Using this new estimate as well as data from previous iterations, a standard (derivative-free) direct search optimization technique computes new values of $\boldsymbol{\theta}$. These new values are passed to the simulation step again and a new iteration starts. The cycle ends when a predefined convergence criterion is satisfied. The final values of the parameters $\boldsymbol{\theta}$ are tested in the evaluation step. Analogously to the simulation step, the resulting control mechanism is applied to n^{eval} demand streams of the evaluation set and various risk measures are calculated. The evaluation step is completely analogous to the simulation step, except that the evaluation set must obviously be independent from the calibration set.

The framework described above can be easily tailored to specific applications by changing two key components that are technically independent from each other: the modified capacity control approach and the optimized risk measure. Accordingly, we will identify the method used with an abbreviation of the form *SBO-[MECHANISM]-[RISKMEASURE]*. In the following, we describe the variants of each of these components we consider in this study. Note that, in addition, the SBO technique can also be varied, but we do not investigate this rather technical issue and stick to a standard approach.

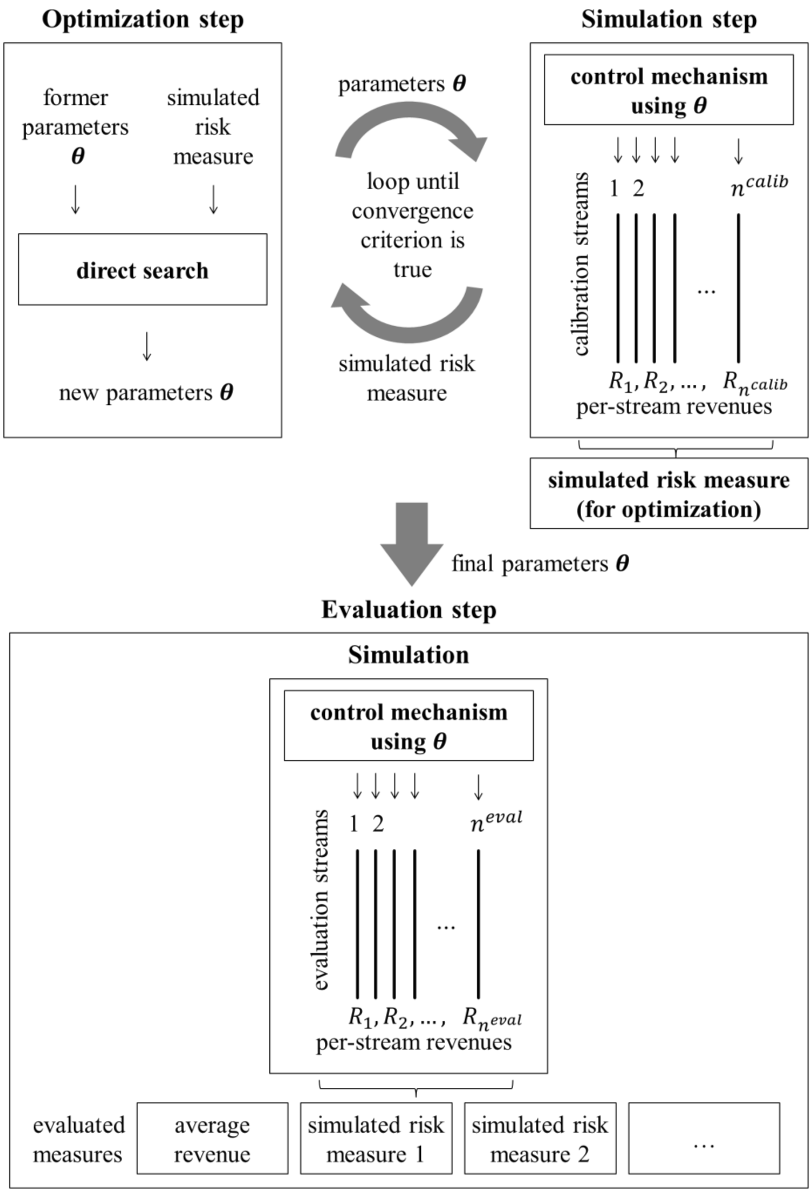

Figure 1: Framework for risk-averse capacity control

Risk measures

In the following, we briefly restate the risk measures used in this study. As customers' arrivals and choices are stochastic, total revenue obtained with a given control mechanism is random and denoted by the random variable R with distribution function $F(y) = \mathbb{P}(R \leq y)$. Note that bigger values of R are preferred.

One well established way to address risk-aversion is the use of expected utility which was introduced by von Neumann and Morgenstern (1944). The main idea behind this concept is that decision makers value the same revenue differently due to individual preferences. These preferences are encompassed in an utility function U and two random revenues, say R_1 and R_2, can be compared by the resulting expected utility, where R_1 is preferred over R_2 if $\mathbb{E}[U(R_1)] \geq \mathbb{E}[U(R_2)]$. Following Barz and Waldmann (2007), we consider an exponential utility function:

$$U_\gamma(R) = 1 - e^{-\gamma \cdot R} \tag{5}$$

The parameter γ indicates the level of risk-aversion. The exponential utility function is the most widely used nonlinear utility function (see, e.g., Corner and Corner (1995)). In our computational study, we abbreviate this risk measure as *Utility*(γ).

The second risk-measure we consider, Conditional Value-at-Risk (CVaR), has attracted a lot of attention over the last decade. For a given probability level $\alpha \in [0,1]$, the CVaR at level α is simply the expectation below the α-quantile of F:

$$CVaR_\alpha(R) = \mathbb{E}[R | R \leq F^{-1}(\alpha)] \tag{6}$$

CVaR is often described as an advancement of the widely popular Value-at-Risk (VaR) to avoid certain theoretical and practical shortcomings of the latter (see, e.g., Artzner et al. (1999)). Note that, to be formally precise, the intuitive definition (6) is valid only for atomless distributions. As revenues are discrete in capacity control, we use CVaR's less intuitive dual representation (not given here; see, e.g., Pflug and Römisch (2007)) to calculate the CVaR. In our computational study, we abbreviate CVaR at level α as *CVaR*(α).

Capacity control mechanisms

In total, we augment three standard control mechanisms for the optimization of arbitrary risk measures.

The first two mechanisms use bid prices π_{tic_i} that, in case of a single resource, come directly from DP-EV (1) or, in case of multiple resources, from the DP decomposition proposed in Liu and van Ryzin (2008). Thus, the bid prices equal or approximate opportunity cost from the risk-neutral problem. Then, building on Koenig and Meissner (2015a) as well as on Huang and Chang (2011), we integrate a constant factor θ_i to adjust the bid prices to different levels of risk-aversion.

Our first mechanism, *BPF* ("**B**id **P**rice control with **F**actor"), follows a traditional, independent demand bid price control approach. Hence, a product j is available for sale if

$$r_j \geq \sum_{i=1}^{m} a_{ij} \cdot \theta_i \cdot \pi_{t+1,ic_i} \tag{7}$$

where the parameters $\boldsymbol{\theta} = (\theta_1, \theta_2, \dots, \theta_m)^T > \mathbf{0}$ are determined using SBO as described above.

In the second mechanism, *AOF* ("**A**ssortment **O**ptimization with **F**actor"), we adjust the bid prices within the exact assortment optimization problem. Compared to (7), this approach is able to consider more combinations of products. Accordingly, the offer set is determined by solving

$$\max_{\boldsymbol{x}} \left\{ \sum_{j=1}^{n} p_{tj}(\boldsymbol{x}) \cdot \left(r_j - \sum_{i=1}^{m} a_{ij} \cdot \theta_i \cdot \pi_{t+1,ic_i} \right) \right\} \tag{8}$$

In order to solve (8) efficiently during our simulation step, we use the greedy algorithm of Miranda Bront et al. (2009). Although this approach is heuristic in nature, it is known to yield high-quality solutions. Please note that the bid prices (i.e., $\theta_i \cdot \pi_{tic_i}$) used in these approaches are artificially set to infinity if $c_i = 0$. Furthermore, they are state-dependent with regard to the state definition from the risk-neutral problem and represent input parameters to the SBO-algorithm altering the bid price control via the choice of $\boldsymbol{\theta}$.

For our third approach, *BPB* ("**B**id **P**rice control with **B**asis functions"), we broadly follow Klein (2007) and use state-dependent bid prices in (7). The state-dependency is given by a linear model of basis functions:

$$\pi_{tic_i} \coloneqq \pi_i^0 - \theta_i^{cap} \cdot \frac{c_i}{c_i^0} + \theta_i^{time} \cdot \frac{(T-t+1)}{T} \tag{9}$$

with $\pi_{tic_i} = \infty$ if $c_i = 0$. Again, the parameters $\boldsymbol{\theta}^{cap} = \left(\theta_1^{cap}, \dots, \theta_m^{cap}\right)^T$ and $\boldsymbol{\theta}^{time} = \left(\theta_1^{time}, \dots, \theta_m^{time}\right)^T$ are estimated by SBO. The variables c_i and $(T-t+1)$ sufficiently describe the current booking situation. Note that we normalize these variables to ease the calibration. π_i^0 is our starting bid price coming from a linear approximation of (DP-EV) such as the well-known deterministic linear program (DLP; see, e.g., Talluri and van Ryzin (1998)) or choice-based deterministic linear program (CDLP; see, e.g., Gallego et al. (2004) and Gallego et al. (2004); Liu and van Ryzin (2008)).

SIMULATION STUDY

In this section, we illustrate the impact of our approaches for a risk-averse decision maker, that is, the improvement in risk measure and, if at all, the associated loss in expected revenue. We use four examples that are—as usual in the literature—expressed in airline terminology. However, the results can be transferred to other areas of application. Wherever available, we use standard example networks from literature.

All algorithms were implemented in MATLAB (Version 8, Release R2014b). Linear Programs were solved by the function linprog from the Optimization Toolbox, Mixed-Integer Linear Programs by CPLEX from IBM ILOG (Version 12.6). In the optimization step, we used the function patternserach with standard settings from the Global Optimization Toolbox. For each problem instance, the size of the evaluation set is $n^{eval} = 10{,}000$. Regarding the three SBO-based approaches presented in the previous section, we use a calibration set of $n^{calib} = 5{,}000$ demand streams. Additional notation introduced in this section is summarized in Table 3.

Single-leg with independent demand		**Parallel flights and one-hub network with choice-based demand**	
p_{tj}	probability of selling product j in period t	l	customer segment
		$\mathcal{C}_l$	consideration set of segment l
Single-leg with choice-based demand		λ_l	arrival probability of a customer from segment l
v_{tj}	preference weight of product j in period t	z_{lj}	binary variable indicating whether consideration set $\mathcal{C}_l$ contains product j
v_0	no-purchase preference weight	$\boldsymbol{v}_l = (v_{lj})_{\lvert\mathcal{C}_l\rvert\times 1}$	preference weights of segment l
		v_{l0}	no-purchase preference weight of segment l

Table 3: Notation introduced in this section

Example 1: Small single-leg flight with independent demand

In our first experiment, we consider the classical single-leg example of Lee and Hersh (1993) which was also used by several previous studies on risk-averse capacity control (see, e.g., Barz (2007), Barz and Waldmann (2007), and Koenig and Meissner (2015a)). It represents a small single-leg flight with a capacity of $c^0 = 10$ seats and $n = 4$ products (booking classes) with revenues $\boldsymbol{r} = (200, 150, 120, 80)^T$. Demand follows the independent demand assumption, that is, the selling probabilities $p_{tj}(\boldsymbol{x})$ are independent of $x_i, i \neq j$, and given by

$$p_{tj}(\boldsymbol{x}) = \begin{cases} p_{tj} & \text{if } x_j = 1 \\ 0 & \text{else} \end{cases} \tag{10}$$

The booking horizon consists of $T = 30$ periods and is partitioned into five time intervals, so that higher value demand tends to arrive later in the booking horizon (see Table 4).

p_{tj}	$j = 1$	$j = 2$	$j = 3$	$j = 4$
$t = 1, \ldots, 5$	0.08	0.08	0.14	0.14
$t = 6, \ldots, 12$	0.06	0.06	0.14	0.14
$t = 13, \ldots, 19$	0.10	0.10	0.10	0.10
$t = 20, \ldots, 26$	0.14	0.14	0.16	0.16
$t = 27, \ldots, 30$	0.15	0.15	0	0

Table 4: Purchase probabilities in Example 1

In this subsection, we consider the risk measures CVaR (6) and expected utility with an exponential utility function (5). We combine these risk measures with the control mechanisms *BPF* and *BPB* and, thus, investigate the performance of *SBO-BPF-CVaR*(α), *SBO-BPB-CVaR*(α), *SBO-BPF-Utility*(γ), and *SBO-BPB-Utility*(γ). Because we assume that demand is independent of the offer set, we do not need to consider the capacity control mechanism *AOF*. Furthermore, we implemented the following approaches as benchmarks:

- BPF_1 is our benchmark. This is the expected revenue-maximizing policy derived from (1), that is, using decision rule (4) with $\pi_{t1c} \coloneqq V_t(c) - V_t(c-1)$.
- $BPF_{0.8}$ uses a constant discount factor of 0.8 on the opportunity cost π_{t1c} in line with Koenig and Meissner (2015a) and Huang and Chang (2011).
- *DP-CVaR*(α) is the CVaR-maximizing policy based on the DP formulation of Gönsch and Hassler (2014) and depends on the probability level α.
- *DP-Utility*(γ) is the expected utility-optimal policy from Barz and Waldmann (2007) and depends on the level of constant absolute risk-aversion γ.

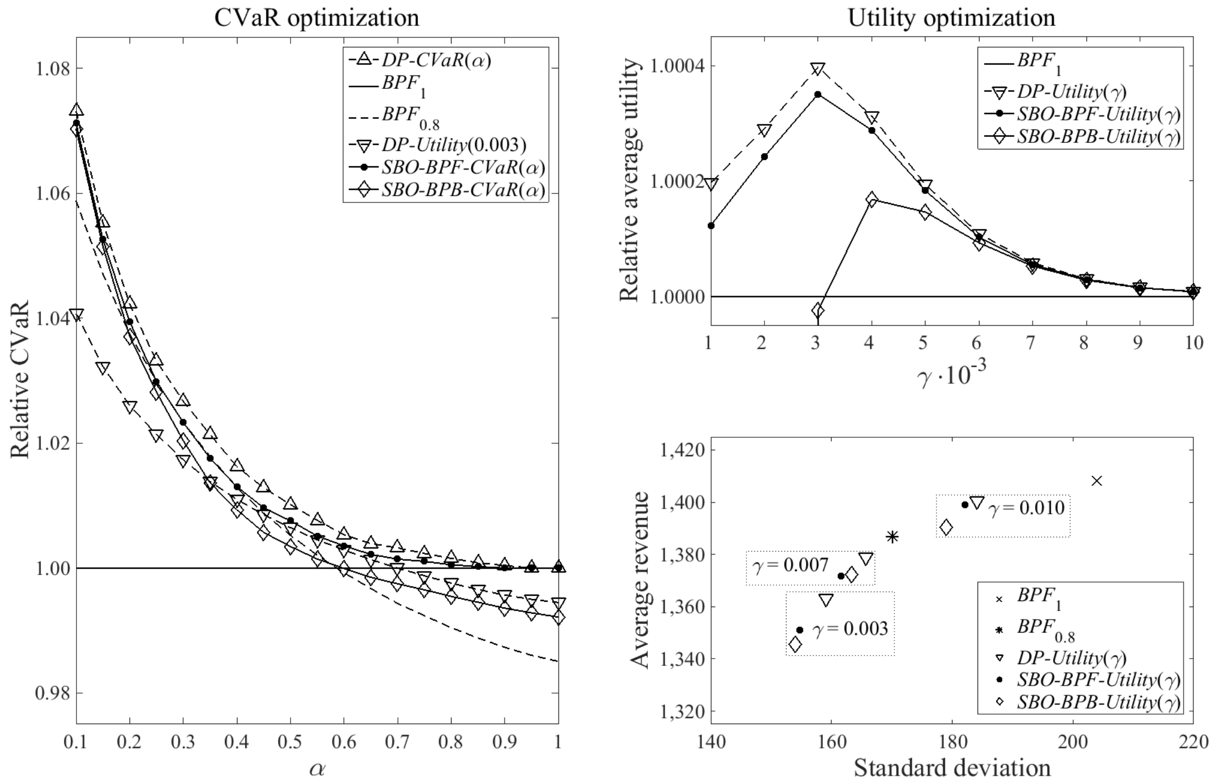

Figure 2: CVaR and average utility in Example 1

In the left (right) part of Figure 2, we consider a CVaR- (utility-) maximizing decision maker and depict the CVaR (utility) relative to that of our benchmark, the expected-value optimal policy from BPF_1. We calculated and evaluated all policies for $\alpha = 0.1, 0.15, \dots, 1$ ($\gamma = 1 \cdot 10^{-3}, \dots, 10 \cdot 10^{-3}$). Taking a look at the left part of Figure 2, not surprisingly, the DP-based approach *DP-CVaR*(α) performs best for all values of α. For $\alpha < 0.6$, all control mechanisms, even *DP-Utility*(0.003), considerably improve CVaR in comparison to BPF_1. A constant discount on opportunity cost ($BPF_{0.8}$), as suggested in previous literature, seems to work very well for $\alpha \in [0.2, 0.4]$ but the results quickly worsen for other values of α. Our simulation-based approach *SBO-BPF-CVaR*(α) is—after *DP-CVaR*(α)—the second best control mechanism for all values of α. The factors θ_1 determined by the SBO monotonically increase from 0.4 to 1 in α. This shows the good performance of the intuitively appealing concept of discounts on opportunity cost, which leads to more accepted requests as risk-aversion increases. Moreover, the fact that *SBO-BPF-CVaR*(α) is able to reclaim most of the difference in CVaR between $BPF_{0.8}$ and *DP-CVaR*(α)—which is applicable only in single-leg settings due to its inherent DP formulation—is encouraging and underlines the performance of the more general SBO approach. However, *SBO-BPB-CVaR*(α) and *DP-Utility*(0.003) yield a poorer performance in this example. This is due to the fact that the linear basis functions of *SBO-BPB-CVaR*(α) are not able to fully capture the monotonicity of the opportunity cost of the expected revenue-maximizing value function (or, equivalently, the concavity of the value function). Therefore, given such a setting, using a simple discount on opportunity cost is advised. Regarding *DP-Utility*(0.003), the poorer performance is not surprising as the corresponding policy is optimized in respect to a different risk measure.

Now, please consider the upper right part of Figure 2. The upper bound on average relative utility is given by the exact DP-based approach *DP-Utility*(γ) of Barz and Waldmann (2007). Because the differences in relative average utility between the different control mechanisms are quite small, we chose to limit the range of values to a small interval, thus excluding $BPF_{0.8}$ from the figure due to a poorer performance. *SBO-BPF-Utility*(γ) works fine for maximizing utility, as the results are nearly identical to *DP-Utility*(γ). The factors θ_1 are again discounts that range from 0.67 to 0.9 and decrease with increasing risk-aversion γ. Similar to the optimization of CVaR, the results of *SBO-BPB-Utility*(γ) are slightly worse. However, all control mechanisms, including BPF_1, show practically identical results for $\gamma \geq 0.007$.

On the lower right part of Figure 2, we compare average revenue and standard deviation for $\gamma \in \{0.003, 0.007, 0.01\}$. Obviously, higher values of γ lead to a smaller average revenue but also a smaller standard deviation of revenues, yielding some kind of efficient frontier. This shows that although the differences in relative utility are often negligible, the approaches lead to different policies.

Example 2: Single-leg flight with choice-based demand

In the remainder of the paper, we assume customer choice behavior. Hence, our main benchmark mechanism is AOF_1 with its near-optimal policy regarding expected revenue and—unless stated otherwise—all results are given relative to this benchmark. Moreover, we now focus on the optimization of CVaR. We consider all three SBO-based mechanisms and additionally state the results of BPF_1. Please note that *DP-CVaR*(α) is not tractable for the following examples.

Unfortunately, we are not aware of an established choice-based single-leg setting from the literature for capacity control. There are only a few settings complementing analytical results. For example, Talluri and van Ryzin (2004a) use a simple example to illustrate demand estimation by an expectation-maximization algorithm and as a proof of concept for their DP formulation. However, in their example, capacity is not scarce (i.e., opportunity cost equals zero) and the authors only have to solve the same assortment optimization problem over time. Nonetheless, the following example is structurally similar.

In this subsection, we consider a single-leg flight with a capacity of $c^0 = 50$, four products with revenues $\boldsymbol{r} = (1000, 800, 600, 400)^T$ and $T = 110$ periods. Demand follows a multinomial logit model. Thus, the purchase probabilities $p_{tj}(\boldsymbol{x})$ depend on product-specific preference weights v_{tj} as well as the no-purchase preference weight $v_0 = 1$ and are given by

$$p_{tj}(\boldsymbol{x}) = \frac{v_{tj} \cdot x_j}{1 + \sum_{k=1}^{n} v_{tk} \cdot x_k} \tag{11}$$

We consider two variants regarding the distribution of demand over time. In the first variant, the purchase probabilities are time-homogenous. In the second variant, higher value demand tends to arrive later in the booking horizon. We call these settings *time-homogenous* and *low-before-high*, respectively. The corresponding values of v_{tj} are given in Table 5. The 10^{-5} values in the second demand variant lead to virtually no demand for the corresponding products, but due to some technicalities, the weights must be strictly positive (also this is often not explicitly stated in the literature).

v_{tj}	$j = 1$	$j = 2$	$j = 3$	$j = 4$
time-homogenous				
$t = 1, \dots, 110$	0.05	0.1	0.5	0.6
low-before-high				
$t = 1, \dots, 80$	10^{-5}	0.2	0.6	0.8
$t = 81, \dots, 110$	0.2	0.3	10^{-5}	10^{-5}

Table 5: Preference weights in Example 2

Figure 3 shows the CVaR of all control mechanisms relative to AOF_1. Note that the spread of relative CVaR is higher in the *low-before-high*-setting. This reflects a well-known effect of capacity control, namely that the decision problem becomes more challenging when demand tends to arrive in low-before-high order.

Naturally, *SBO-AOF-CVaR*(α) performs best and, in both examples, with small benefits over *SBO-BPB-CVaR*(α) as well as larger benefits over *SBO-BPF-CVaR*(α). Interestingly, even the standard bid price control BPF_1 is often competitive and yields a higher CVaR than AOF_1 for low levels of α because more low value products are offered for sale. However, for medium and high values of α, there are severe losses in CVaR. With SBO, these losses can be successfully reduced.

Regarding θ_1, we observe values of 0.45 to 1 for *SBO-AOF-CVaR*(α) that are increasing in α and which represent discounts on the opportunity cost analogously to Example 1. Regarding *SBO-BPF-CVaR*(α), θ_1 ranges from 0.85 to 1.87 and also increases with α. This reflects that with a bid price control, there is a trade-off between a discount on the opportunity cost to allow for risk-aversion and a markup to prohibit buy-down behavior. This effect is better captured by the state-dependent bid prices of *SBO-BPB-CVaR*(α) in comparison with *SBO-BPF-CVaR*(α), which uses a constant markup over the whole booking horizon. In a risk-neutral setting, increased bid prices to induce higher value demand were, for example, also observed in Meissner and Strauss (2012a).

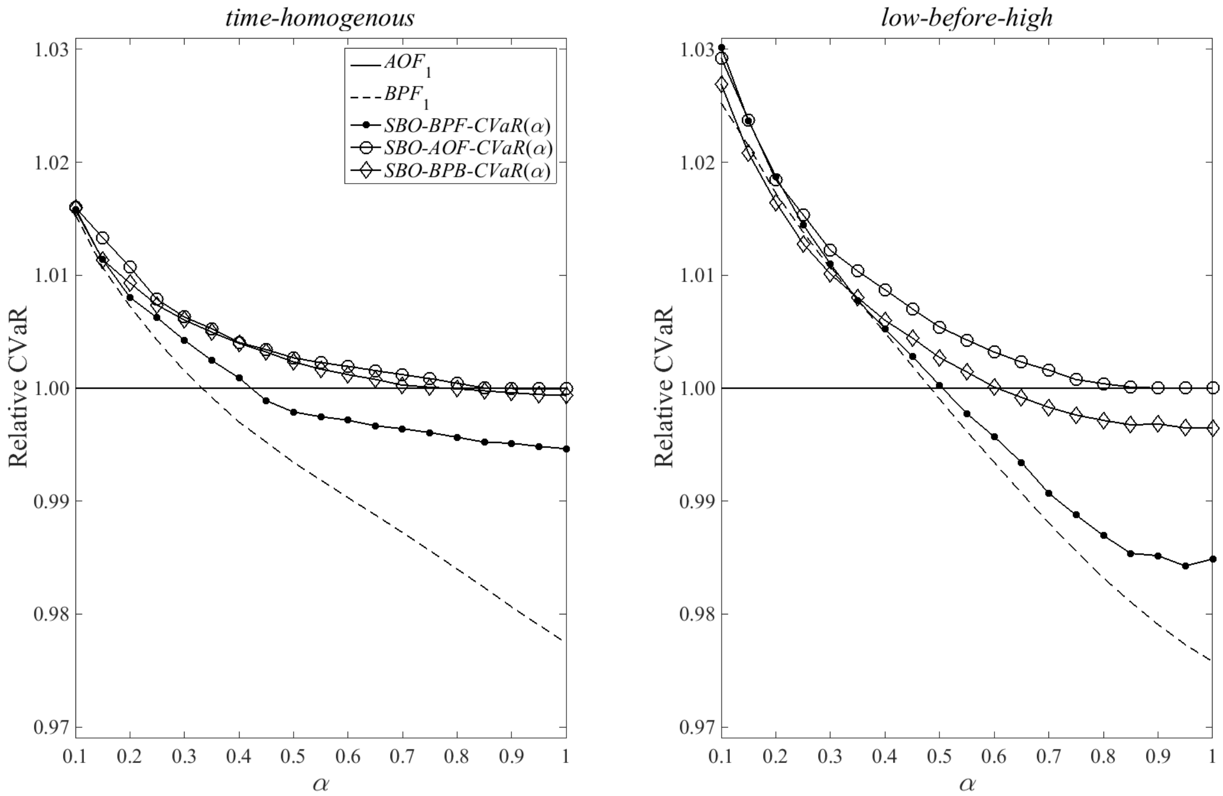

Figure 3: CVaR in Example 2

We now delve deeper into the opportunities and threats that accompany risk-averse capacity control. For all SBO-based mechanisms and three selected values of α, the upper part of Figure 4 shows the absolute gains in CVaR compared to the absolute gains in average revenue. Please note that the gains in CVaR are subject to the specific level of risk-aversion α and, thus, need to be treated with caution. A higher gain in CVaR usually leads to a bigger loss in expected revenue. For example, in *low-before-high*, improving the $\mathrm{CVaR}_{0.4}$ by around 250 costs 500 in average revenue (see the upper right part of Figure 4). The lower part of Figure 4 compares the sampled distribution of total revenues of *SBO-AOF-CVaR*(0.4) and AOF_1 over the evaluation streams. In line with the results from the optimization of utility in Example 1, risk-averse capacity controls leads to a smaller support of the distribution and shorter tails. In other words, extreme outcomes are less likely.

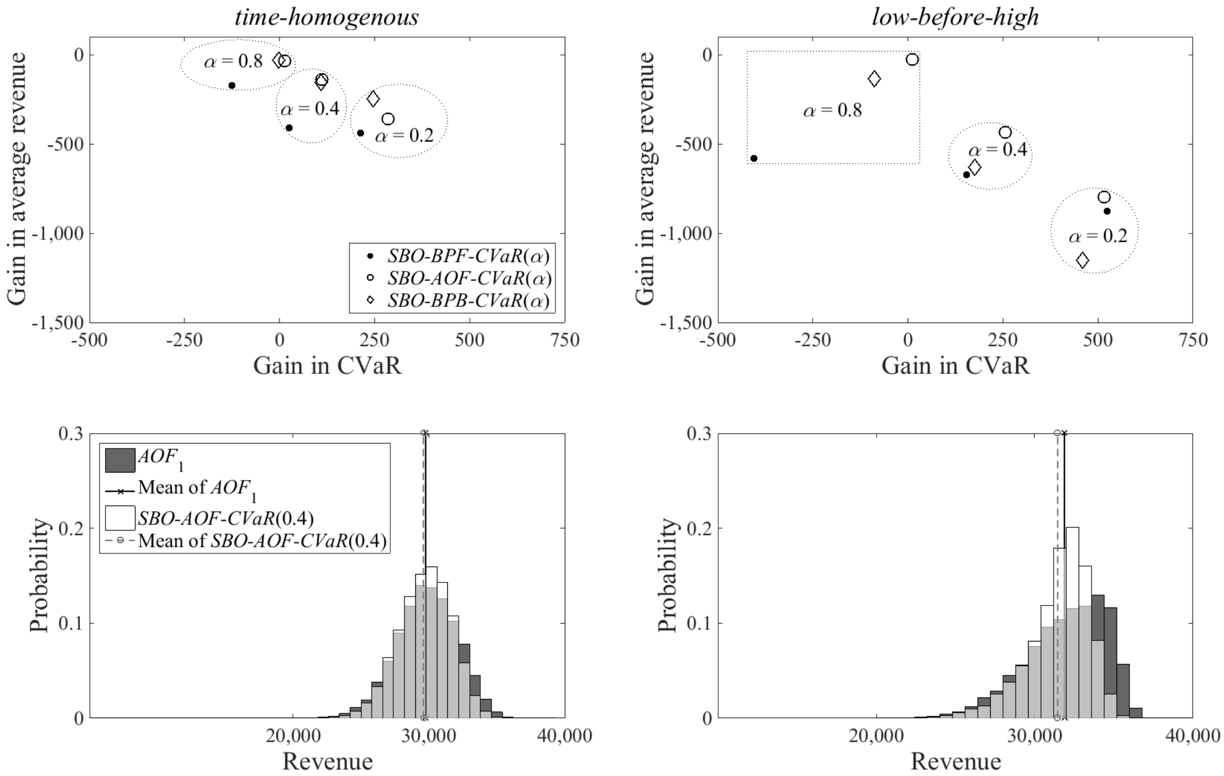

Figure 4: CVaR vs. expected value and revenue distribution in Example 2

Example 3: Parallel flights with choice-based demand

Our third example is based on the parallel flight network of Miranda Bront et al. (2009). It consists of three flights with two products defined on each flight, that is, one low-class and one high-class product. In this example, we additionally investigate the impact of different capacity provision on the risk-profile: First, we consider an initial capacity of $\boldsymbol{c}^0 = (27, 45, 36)^T$ and, second, an initial capacity of $\boldsymbol{c}^0 = (21, 35, 28)^T$. Demand follows a mixture of multinomial logit models. More precisely, customers belong to different market segments $l = 1, \ldots, 4$, each of which has a subset of products to consider for purchase, namely the consideration set $\mathcal{C}_l$. The variable z_{lj} indicates whether product $j \in \mathcal{C}_l$ ($z_{lj} = 1$) or not ($z_{lj} = 0$). A customer from segment l arrives with probability λ_l and has preference weights $\boldsymbol{v}_l = \left(v_{lj}\right)_{|\mathcal{C}_l| \times 1}$ as well as v_{l0} for the no-purchase alternative. Note that v_{lj} is only defined if $z_{lj} = 1$. Demand is time-homogenous over the booking horizon of $T = 300$ periods. Then, the probability of selling product j in period t is given by

$$p_{tj}(\boldsymbol{x}) = \sum_{l=1}^{4} \lambda_l \cdot \frac{v_{lj} \cdot z_{lj} \cdot x_j}{v_{l0} + \sum_{k=1}^{n} v_{lk} \cdot z_{lk} \cdot x_k} \qquad (12)$$

The remaining data is summarized in Table 6 in Appendix A. The corresponding relative CVaRs are shown in Figure 5. Similar to the previous example, the absolute CVaR gains vs.

absolute expected revenue gains and the sampled distribution of total revenues for $\alpha = 0.4$ are illustrated in Figure 6.

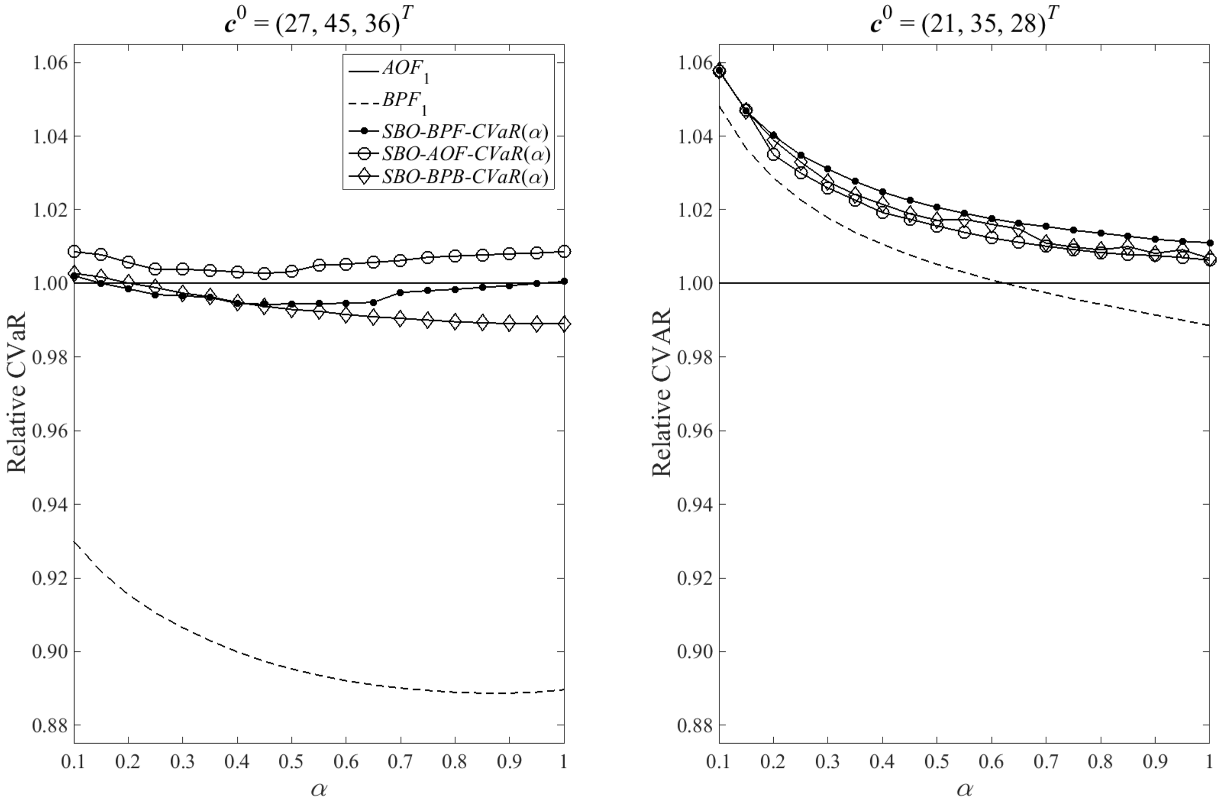

Figure 5: CVaR in Example 3

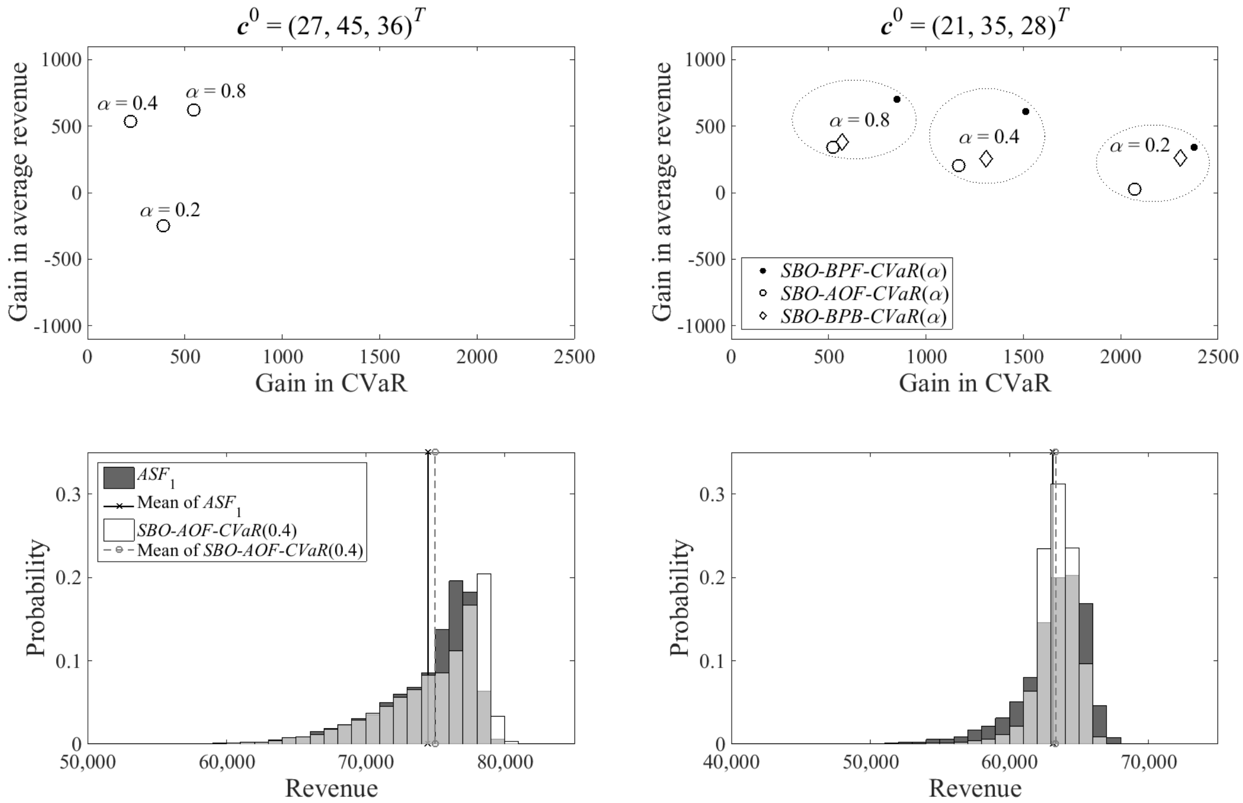

Figure 6: CVaR vs. expected value and revenue distribution in Example 3

In the following, we summarize the key observations complementing the former results. First, comparing the results for the two initial capacities, we see that only minor variations of the setting can lead to large differences in the risk profile. Second, there are settings, such as the first initial capacity, where the consideration of risk-aversion is more or less negligible. Given such a setting, the better the expected revenue of a policy is, the better is its CVaR for almost all levels α and vice versa. More precisely, optimizing one can also increase the other and it suffices to optimize expected revenue. Third, standard bid price controls such as BPF_1 can perform quite poorly when considering customer choice behavior (second initial capacity) and SBO can successfully address this. In this instance, *SBO-BPF-CVaR*(α) uses factors θ_i of up to 3.97 and outperforms the other approaches. It is usually even better than the near-exact assortment optimization in AOF_1. For example, the gain of *SBO-AOF-CVaR*(1) in expected revenue over AOF_1 is almost 1%. This remarkable result can only be explained with the fact that all approaches, including AOF_1, use approximate bid prices from a DP decomposition instead of the exact opportunity cost from the intractable DP formulation.

Example 4: One hub-network with choice-based demand

The last example is based on an airline network structure from Meissner and Strauss (2012b) with one hub H connecting two non-hub cities A and B with four flight legs (see Figure 7). There are six itineraries (A to H, A to B via H, H to B, B to H, B to A via H, and H to A). For each itinerary, one high-class and one low-class product are available. The demand behavior is the same as in the parallel flight example. For each itinerary, there is one customer segment with a higher preference for the low-class product. A detailed description of products (revenues r_j and capacity consumptions $\boldsymbol{a}_j$) and segments (consideration sets $\mathcal{C}_l$, preference weights $\boldsymbol{v}_l$, no-purchase preference weights v_{l0}, and segment probabilities λ_l) can be found in Table 7 in Appendix A.

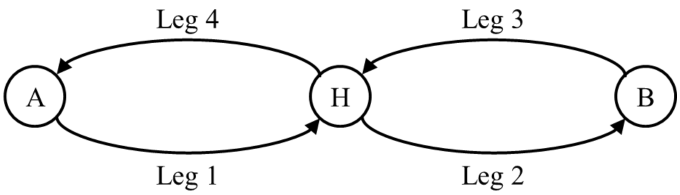

Figure 7: One-hub network of Example 4

We assume an initial capacity $c_h^0 = 15$ ($c_h^0 = 60$) for all h and a booking horizon length $T = 300$ ($T = 1200$). The corresponding relative CVaRs are shown in Figure 8.

In general, the observations confirm the results of the previous examples. However, it is remarkable that large gains in CVaR and expected revenue in settings with connecting flights are possible. This is because AOF_1 decomposes the network by flights and, apparently, this does not sufficiently capture the network effects. SBO is able to remedy this shortcoming and, thus, all three SBO-based approaches perform quite well. Comparing the two initial capacities shows that the differences between all mechanisms decline as we scale up the size of the network.

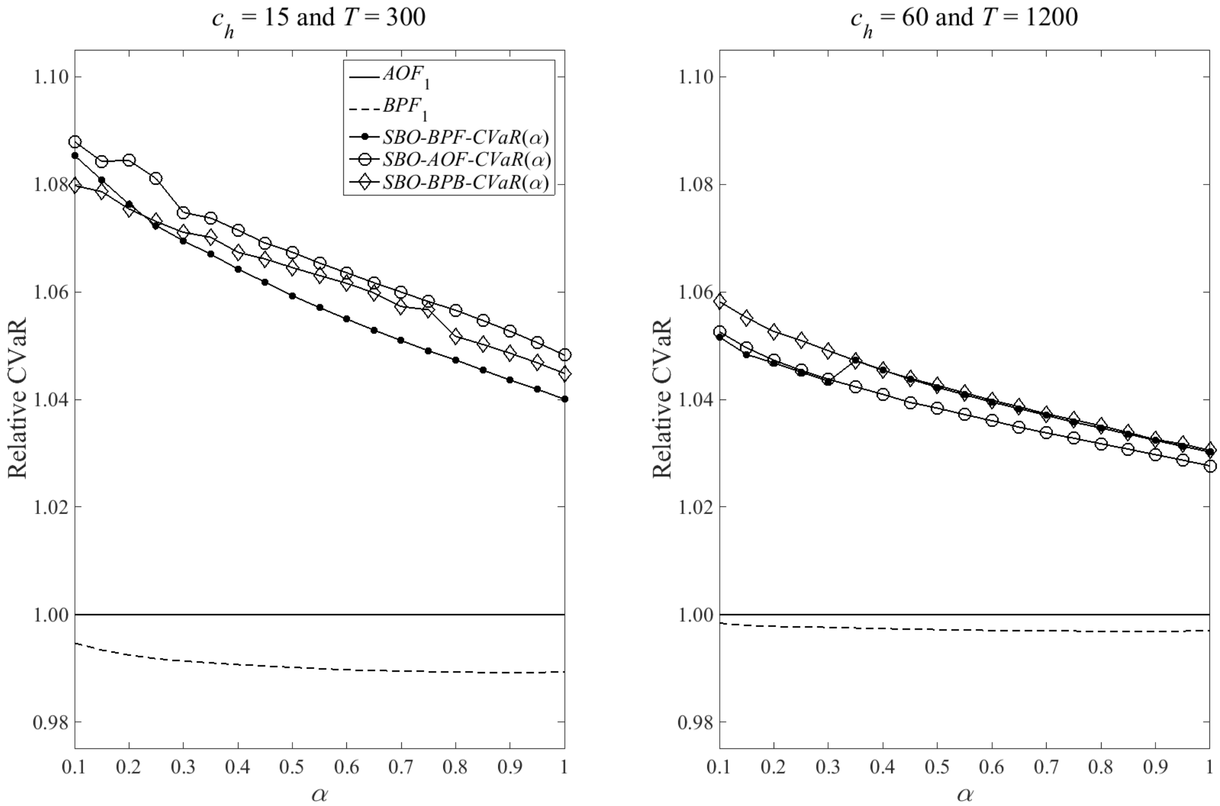

Figure 8: CVaR in Example 4

DISCUSSION

After analyzing the numerical results in detail in the previous section, we now take a broader perspective and discuss the relevance of our work.

First, the relevance of risk-aversion in revenue management—and particularly network revenue management—is the foundation of this paper, although risk-neutrality has been taken for granted for a long time in the literature. One reason is probably that risk-neutrality leads to mathematically simpler models. However, many people who are new to revenue management consider this assumption counter-intuitive and many industry partners question it at first.

Lancaster (2003) was the first to raise the issue of risk in revenue management. He pointed out that airlines, like all businesses, face risks which should be managed appropriately. By contrast, he observes that revenue management considers only the reward side, that is, increasing expected revenue and completely ignores the risks assumed in doing so. Despite this early work, most authors cite experiences from practice to show the relevance of risk-aversion. For example, smaller airlines asked a consultant about risk-averse capacity control (see Weatherford (2004)). Two other examples are due to Levin et al. (2008). Event promoters may organize only a few large events per year in locations that are very expensive to rent. Accordingly, their first priority is to recover this fixed cost (see Levin et al. (2008)). In other industries, a manager's primary concern is often to provide stable results because negative news can lead to negative stock market assessments that can far outweigh the marginal revenue advantages of a risk-neutral policy.

In contrast to the above-mentioned rather small businesses, many companies have a large number of similar events. Thus, the law of large numbers ensures that the average revenue of each event is maximized and also quite stable when using a risk-neutral model that focuses on the expected revenue. For example, network airlines have hundreds, major ones even several thousands of take-offs every day. Although risk-neutrality may be appropriate for these companies as a whole, it may not be appropriate for every department and individual decision maker, leading to missing acceptance of risk-neutral revenue management systems. For example, a consultant's clients were not comfortable with their risk-neutral revenue management system (see Barz (2007)). They manually altered the forecast to obtain less aggressive (and risky) results. Singh (2011) observed that analysts' individual risk-aversion has a huge impact on their decisions when overwriting a revenue management system's output at a cruise line company. He attributes this behavior largely to their personality because they made decisions about exactly the same issues and possessed identical information.

Second, we would like to point out that, regarding the company as a whole, the need to incorporate risk-aversion declines with the network size due to compensatory effects. However, our examples showed that there are also bigger settings where risk-aversion is relevant in the sense that a risk-averse solution differs from a risk-neutral one. For this purpose, we used established examples from literature both for single-leg and network capacity control with mid-sized capacity. In our opinion, these networks serve as a good representation of the sub-networks that individual decision makers may control. Nonetheless, it would be worthy of future research to investigate the impact of risk-averse control mechanisms in large-scale problem instances from practice.

Third, our results show that we were able to sufficiently address customer choice behavior in most cases by using bid price rules instead of solving the exact assortment optimization problem. This is in line with the literature on the optimization of expected revenue: Chaneton and Vulcano (2011) and Meissner and Strauss (2012a) make similar observations. But, despite of their widespread use, bid price controls can sometimes yield a comparatively poor performance. This may be due to the fact that bid price controls are not always able to represent the optimal policy in networks, especially if customer choice is considered (see, e.g., Talluri and van Ryzin (1998)). By contrast, the solution of the assortment optimization problem is able to represent all decision options. Thus, *SBO-AOF-CVaR*(α) performs slightly better than the bid price controls. Nonetheless, our results indicate that *SBP-BPF-CVaR*(α) and *SBO-BPB-CVaR*(α) perform very well and, thus, explain the favoritism of bid price controls in practice due to the trade-off between solution quality and simplicity.

To summarize the discussion so far, the overall framework works quite well because arbitrary controls may be designed and optimized regarding a certain risk measure. We focused on enhancing existing control mechanisms with risk-averse components, since commercial revenue management systems are fixed in the long run. Given a control mechanism, our framework always improves the results of the original control.

Finally, the framework presented can also be used to capture risk-aversion in dynamic pricing, where a firm decides on the products' prices instead of their availability. Thus, instead of the assortment optimization problem (3), the firm has to solve a pricing problem in each period to determine the products' prices. As in capacity control, the value of future sales is reflected by opportunity cost which is often approximated by bid prices. Therefore, our framework can be easily applied to risk-averse dynamic pricing because the bid prices can be modified via tunable parameters that capture risk-aversion. Moreover, if continuous prices are allowed, the pricing problem can be solved analytically for many demand models and, thus, much faster than the assortment optimization problem we consider. This could allow the SBO to perform more simulation runs, possibly leading to even better results.

CONCLUSION

We presented a flexible and modular framework for risk-averse capacity control that offers several advantages compared to existing approaches. First, the practical decision rules we consider can be implemented easily in existing operational systems because they build on well-established standard risk-neutral control mechanisms. Second, using SBO, the control mechanisms can be tailored to every risk measure. Third, because SBO-algorithms are meanwhile widely available in standard software, the only prerequisite for using this model-free framework is being able to undertake Monte-Carlo simulations of the arrival process and choice behavior of customers. There is no need for a DP formulation of the decision problem, which is prohibitive for most risk measures, or an explicit model of customer behavior.

For demonstration purposes, we conducted a simulation study with the widely used multinomial logit model, but our approach admits the use of any other choice model. Based on CVaR and expected utility, we showed that small and intuitive modifications in standard control mechanisms, if designed properly, can be sufficient to successfully incorporate risk-aversion into capacity control, including network settings and customer choice. This usually leads to a narrower distribution of revenues in comparison to standard controls as well as more predictable and stable revenues.

Appendix

Appendix A

Regarding Example 3 (parallel flights) and Example 4 (one hub), Table 6 and Table 7 summarize the remaining product and segment data.

Products			Segments				
j	$\boldsymbol{a}_j$	r_j	l	$\mathcal{C}_l$	$\boldsymbol{v}_l$	v_{l0}	λ_l
1	$(1,0,0)^T$	400	1	$\{2,4,6\}$	$(5,10,1)^T$	1	0.1
2	$(1,0,0)^T$	800	2	$\{1,3,5\}$	$(5,1,10)^T$	5	0.15
3	$(0,1,0)^T$	500	3	$\{1,2,3,4,5,6\}$	$(10,8,6,4,3,1)^T$	5	0.2
4	$(0,1,0)^T$	1000	4	$\{1,2,3,4,5,6\}$	$(8,10,4,6,1,3)^T$	1	0.05
5	$(0,0,1)^T$	300					
6	$(0,0,1)^T$	600					

Table 6: Product and segment description in Example 3

Products			Segments				
j	$\boldsymbol{a}_j$	r_j	l	$\mathcal{C}_l$	$\boldsymbol{v}_l$	v_{l0}	λ_l
1	$(1,0,0,0)^T$	300	1	$\{1,2\}$	$(0.5,2)^T$	1	0.1
2	$(1,0,0,0)^T$	150	2	$\{3,4\}$	$(0.5,2)^T$	1	0.06
3	$(1,1,0,0)^T$	600	3	$\{5,6\}$	$(0.5,2)^T$	1	0.1
4	$(1,1,0,0)^T$	300	4	$\{7,8\}$	$(0.5,2)^T$	1	0.1
5	$(0,1,0,0)^T$	350	5	$\{9,10\}$	$(0.5,2)^T$	1	0.09
6	$(0,1,0,0)^T$	175	6	$\{11,12\}$	$(0.5,2)^T$	1	0.07
7	$(0,0,1,0)^T$	300					
8	$(0,0,1,0)^T$	150					
9	$(0,0,1,1)^T$	500					
10	$(0,0,1,1)^T$	250					
11	$(0,0,0,1)^T$	250					
12	$(0,0,0,1)^T$	125					

Table 7: Product and segment description in Example 4

References

Artzner, P., Delbaen, F., Eber, J. and Heath, D. (1999) Coherent measures of risk. *Mathematical Finance* 9(3): 203–228.

Barz, C. (2007) *Risk-Averse Capacity Control in Revenue Management.* Berlin, Heidelberg: Springer.

Barz, C. and Waldmann, K.H. (2007) Risk-sensitive capacity control in revenue management. *Mathematical Methods of Operations Research* 65(3): 565–579.

Belobaba, P. (1992) Optimal vs. heuristic methods for nested seat allocation. Presentation at AGIFORS Reservations and Yield Management Study Group, Brussels.

Bertsimas, D. and de Boer, S. (2005) Simulation-based booking limits for airline revenue management. *Operations Research* 53(1): 90–106.

Chaneton, J.M. and Vulcano, G. (2011) Computing bid prices for revenue management under customer choice behavior. *Manufacturing & Service Operations Management* 13(4): 452–470.

Corner, J.L. and Corner, P.D. (1995) Characteristics of decisions in decision analysis practice. *Journal of the Operational Research Society* 46(3): 304–314.

Feng, Y. and Xiao, B. (1999) Maximizing revenues of perishable assets with a risk factor. *Operations Research* 47(2): 337–341.

Feng, Y. and Xiao, B. (2008) A risk-sensitive model for managing perishable products. *Operations Research* 56(5): 1305–1311.

Gallego, G., Iyengar, G., Phillips, R. and Dubey, A. (2004) Managing Flexible Products On A Network. Working paper, Columbia University, New York.

Gallego, G. and van Ryzin, G. (1997) A multiproduct dynamic pricing problem and its applications to network yield management. *Operations Research* 45(1): 24–41.

Gönsch, J. and Hassler, M. (2014) Optimizing the conditional value-at-risk in revenue management. *Review of Managerial Science* 8(4): 495–521.

Gönsch, J., Hassler, M. and Schur, R. (2015) Optimizing conditional value-at-risk in dynamic pricing. Working paper, University of Augsburg, Augsburg.

Gosavi, A. (2015) *Simulation-Based Optimization: Parametric Optimization Techniques and Reinforcement Learning,* 2nd ed., New York: Springer.

Gosavi, A., Ozkaya, E. and Kahraman, A.F. (2007) Simulation optimization for revenue management of airlines with cancellations and overbooking. *OR Spectrum* 29(1): 21–38.

Huang, K. and Chang, K.-C. (2011) A model for airline seat control considering revenue uncertainty and risk. *Journal of Revenue and Pricing Management* 10(2): 161–171.

Klein, R. (2007) Network capacity control using self-adjusting bid-prices. *OR Spectrum* 29(1): 39–60.

Koenig, M. and Meissner, J. (2015a) Risk management policies for dynamic capacity control. *Computers & Operations Research* 59: 104–118.

Koenig, M. and Meissner, J. (2015b) Risk minimising strategies for revenue management problems with target values. *Appears in: Journal of the Operational Research Society*.

Koenig, M. and Meissner, J. (2015c) Value-at-risk optimal policies for revenue management problems. *International Journal of Production Economics* 166: 11–19.

Lancaster, J. (2003) The financial risk of airline revenue management. *Journal of Revenue and Pricing Management* 2(2): 158–165.

Lee, T.C. and Hersh, M. (1993) A model for dynamic airline seat inventory control with multiple seat bookings. *Transportation Science* 27(3): 252–265.

Levin, Y., McGill, J. and Nediak, M. (2008) Risk in revenue management and dynamic pricing. *Operations Research* 56(2): 326–343.

Li, M.Z. and Zhuang, W. (2009) Risk-sensitive dynamic pricing for a single perishable product. *Operations Research Letters* 37(5): 327–332.

Liu, Q. and van Ryzin, G. (2008) On the choice-based linear programming model for network revenue management. *Manufacturing & Service Operations Management* 10(2): 288–310.

Meissner, J. and Strauss, A. (2012a) Improved bid prices for choice-based network revenue management. *European Journal of Operational Research* 217(2): 417–427.

Meissner, J. and Strauss, A. (2012b) Network revenue management with inventory-sensitive bid prices and customer choice. *European Journal of Operational Research* 216(2): 459–468.

Miranda Bront, J.J., Méndez-Díaz, I. and Vulcano, G. (2009) A column generation algorithm for choice-based network revenue management. *Operations Research* 57(3): 769–784.

Pflug, G. and Römisch, W. (2007) *Modeling, Measuring and Managing Risk.* Singapore: World Scientific.

Phillips, R.L. (2005) *Pricing and Revenue Optimization.* Stanford: Stanford University Press.

Robinson, L.W. (1995) Optimal and approximate control policies for airline booking with sequential nonmonotonic fare classes. *Operations Research* 43(2): 252–263.

Singh, P. (2011) Risk in revenue management. Acknowledging risk's existence and knowing how to minimize it. *Informs Analytics Magazine* May/June: 16–19.

Spall, J.C. (2003) *Introduction to Stochastic Search and Optimization: Estimation, Simulation, and Control.* Hoboken, New Jersey: John Wiley & Sons.

Talluri, K. and van Ryzin, G. (1998) An analysis of bid-price controls for network revenue management. *Management Science* 44(11, part 1): 1577–1593.

Talluri, K.T. and van Ryzin, G. (2004a) Revenue management under a general discrete choice model of consumer behavior. *Management Science* 50(1): 15–33.

Talluri, K.T. and van Ryzin, G. (2004b) *The Theory and Practice of Revenue Management.* New York: Springer.

Topaloglu, H. (2008) A stochastic approximation method to compute bid prices in network revenue management problems. *INFORMS Journal on Computing* 20(4): 596–610.

van Ryzin, G. and Vulcano, G. (2008a) Computing virtual nesting controls for network revenue management under customer choice behavior. *Manufacturing & Service Operations Management* 10(3): 448–467.

van Ryzin, G. and Vulcano, G. (2008b) Simulation-based optimization of virtual nesting controls for network revenue management. *Operations Research* 56(4): 865–880.

von Neumann, J. and Morgenstern, O. (1944) *Theory of Games and Economic Behavior.* Princeton: Princeton University Press.

Weatherford, L.R. (2004) EMSR versus EMSU: Revenue or utility? *Journal of Revenue and Pricing Management* 3(3): 277–284.

Zhang, D. and Adelman, D. (2009) An approximate dynamic programming approach to network revenue management with customer choice. *Transportation Science* 43(3): 381–394.

Zhuang, W. and Li, M.Z.F. (2011) Revenue optimization of risk-averse managers with atemporal utility. *Journal of Revenue and Pricing Management* 10(5): 424–437.