

Optimizing the Conditional Value-at-Risk in Revenue Management

July 10, 2013

Jochen Gönsch, Michael Hassler

Jochen Gönsch (✉)

*Department of Analytics & Optimization,
University of Augsburg, Universitätsstraße 16, 86159 Augsburg, Germany*

Phone: +49 821 598 4148

Fax: +49 821 598 4226

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

Michael Hassler

*Department of Analytics & Optimization,
University of Augsburg, Universitätsstraße 16, 86159 Augsburg, Germany*

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

Abstract

Many service industries use revenue management to balance demand and capacity. The assumption of risk-neutrality lies at the heart of the classical approaches, which aim at maximizing expected revenue. In this paper, we give a comprehensive overview of the existing approaches, most of which were only recently developed, and discuss the need to take risk-averse decision makers into account. We then present a heuristic that maximizes Conditional Value-at-Risk (CVaR). Although CVaR has become increasingly popular in finance and actuarial science due to its beneficial properties, this risk measure has not yet been considered in the context of revenue management. We are able to efficiently solve the optimization problem inherent in CVaR by taking advantage of specific structural properties that allow us to reformulate this optimization problem as a continuous knapsack problem. In order to demonstrate the applicability and robustness of our approach, we conduct a simulation study that shows that the new approach can significantly improve the risk profile in various scenarios.

*Keywords: Revenue Management, Capacity Control, Dynamic Decisions,
Dynamic Programming, Risk-Aversion, Conditional Value-at-Risk*

JEL Classification M11, M19, C61, C63

1 Introduction

Revenue management is “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). It combines results from areas like demand modeling and forecasting as well as mathematical optimization with the goal of maximizing expected revenue (or profit). The resulting models and procedures are usually integrated into automated systems seeking to manage demand by restricting the availability of products through some kind of capacity control.

All the classical approaches have in common that the assumption of a risk-neutral decision maker lies at their heart. This was justified by the large number of similar decisions in typical fields of application like airlines and hotel chains. Although many people who are new to revenue management consider this assumption counter-intuitive and many industry partners question it at first, it has been taken for granted for a long time. One reason is that it leads to mathematically simpler models, maximizing expected values.

The contribution of this paper is twofold: On the one hand, we introduce the topic and provide the first comprehensive, up-to-date overview of existing approaches that consider risk-aversion in revenue management’s capacity control. On the other hand, we present a new approach that – to the best of our knowledge – is the first one aiming at maximizing Conditional Value-at-Risk (CVaR) and conduct a simulation study, showing how the heuristic improves the risk profile of the resulting policy in different scenarios.

The remainder of this paper is organized as follows: In Section 2, we give an introduction to revenue management and point out some further literature to the reader. Section 3 discusses when the assumption of risk-neutrality is and is not appropriate and provides a comprehensive overview of the existing approaches to incorporating risk. The new approach that optimizes CVaR is presented in Section 4. The design of our simulation study is described in Section 5 and the results are discussed in Section 6. We conclude with a summary and an outlook on avenues of further research in Section 7.

2 Revenue Management

The deregulation of the US airline industry in the 1970s gave rise to revenue management as a number of new competitors offering cheap ticket prices entered the market. Whereas business customers largely remained loyal to the established airlines because of their extensive flight networks with carefully orchestrated connection flights and high frequencies, leisure travelers switched to the new low-cost airlines wherever they offered a point-to-point service between two cherry-picked cities. This left the traditional airlines in a dilemma: to stay profitable, they had to regain leisure travelers, but because of the high cost resulting from operating a complex flight network, they could not lower their prices to the low-cost airlines' price level. American Airlines was the first to solve this problem with the introduction of an additional "Super Saver Fare" aimed solely at leisure travelers. This fare was subjected to purchase restrictions, preventing business travelers from switching to the new, competitively priced tickets: for example, tickets had to be purchased 30 days in advance and required a minimum stay of seven days. However, the popularity of these new fares with leisure customers raised a new question. Clearly, the new approach would not be beneficial if the sale of a low-price ticket replaced a high-paying business customer. Thus, how many seats on each flight should be sold at the "Super Saver Fare"? This question led to the development of revenue management.

Despite its birth in practice, revenue management has been extensively researched over the past 20 years, as, for example, the surveys by Weatherford and Bodily (1992), McGill and van Ryzin (1999), Tscheulin and Lindenmeier (2003) or Chiang et al. (2007) show. Furthermore, the standard textbooks by Talluri and van Ryzin (2004) and Phillips (2005) give an overview of the field as well as an in-depth discussion.

Revenue management instantly became a must-have in the airline industry (see Cross 1997). It spread to other service industries where firms have a largely fixed capacity and face stochastic demand of inhomogeneous value. Examples are car rental companies (see, e.g., Carrol and Grimes 1995, Geraghty and Johnson

1997), tour operators (see, e.g., Hoseason and Johns 1998, Klein 2000, Xylander 2003), passenger railways (see, e.g., Ciancimino et al. 1999, Hood 2000, Bharill and Rangaraj 2008, You 2008) to name but a few. More recently, revenue management has also been considered in manufacturing (see, e.g., Barut and Sridharan 2005, Rehkopf 2006, Spengler et al. 2007, Defregger and Kuhn 2007, Wiggershaus 2008, Hintsches et al. 2010, and Volling et al. 2012).

Following Lautenbacher and Stidham (1999), the standard revenue management problem of capacity control is usually formulated as a markov decision process (see also the abovementioned surveys and textbooks). In this problem, a firm disposes of a perishable resource with a capacity of C units. Products $i = 1, \dots, n$ each make use of one unit of this capacity and are sold during a selling horizon of finite length T . Any capacity that remains thereafter is worthless. Without loss of generality, the products are priced at $0 \leq r_n \leq r_{n-1} \leq \dots \leq r_1$ and variable costs are generally neglected (although they could be integrated straightforwardly). Regarding demand, the selling horizon is divided into a sufficiently large number of T micro periods such that, in each micro period, no more than one customer requests a product. Time is indexed forwards, so micro periods 1 and T mark the beginning respectively the end of the selling horizon. Furthermore, the common independent demand (ID) assumption holds; that is, the probability $p_i(t)$ that product i is requested in micro period $t = 1, \dots, T$ is given ex ante and $p_0(t) = 1 - \sum_{i=1}^n p_i(t)$ denotes the probability that no product is requested. The optimal expected revenue $V_t^N(c)$ from period t onwards with remaining capacity c is now given by the Bellman equation

$$V_t^N(c) = \sum_{i=1}^n p_i(t) \cdot \max_{x_{ti} \in \{0,1\}} \{x_{ti} \cdot r_i + V_{t+1}^N(c - x_{ti})\} + p_0(t) \cdot V_{t+1}^N(c) = \max_{x_t \in \{0,1\}^n} \sum_{i=1}^n p_i(t) \cdot (x_{ti} \cdot r_i + V_{t+1}^N(c - x_{ti})) + p_0(t) \cdot V_{t+1}^N(c) \quad (1)$$

with the boundary conditions $V_t^N(c) = -\infty$ for $c < 0$ and $t = 1, \dots, T$ and $V_T^N(c) = 0$ for $c \geq 0$. The variable $x_{ti} \in \{0,1\}$ describes the denial ($x_{ti} = 0$) or acceptance ($x_{ti} = 1$) of a request for product i . This formulation can be explained

as follows: In each micro period, a request for product i arrives with probability $p_i(t)$. If a request for product i arrives, the firm can either accept or reject it. If it is accepted ($x_{ti} = 1$), the firm immediately collects revenue r_i , but goes into the next period with one fewer unit of capacity, where it obtains an expected revenue of $V_{t+1}^N(c - 1)$. If the request is rejected ($x_{ti} = 0$), or no request arrives, the firm expects a future revenue of $V_{t+1}^N(c)$, calculated with an unchanged capacity. Thus, we obtain the optimal expected revenue with full capacity C over the entire selling horizon by recursively calculating $V_1^N(C)$.

It is now easy to see that our firm should accept a request in micro period t if and only if $r_i + V_{t+1}^N(c - 1) \geq V_{t+1}^N(c)$ or, equivalently,

$$r_i \geq V_{t+1}^N(c) - V_{t+1}^N(c - 1). \quad (2)$$

Here, the right hand side is the opportunity cost or marginal value of capacity and (2) states that a request is accepted only if the associated revenue exceeds the opportunity cost. By calculating (1) for every $t = 1, \dots, T$ and $c = 1, \dots, C$, the firm obtains a decision rule (policy) that guides decision making in every state (c, t) in order to maximize the expected revenue.

3 Considering Risk in Revenue Management

In this section, we review the consideration of risk in the revenue management literature. For an overview of research incorporating risk in adjacent areas like newsvendor problems or dynamic pricing, see, for example, Barz (2007, Chapter 1.2) and Gönsch et al. (2013, Chapter 4.8).

In revenue management, the consideration of expected values is generally justified by the large number of similar selling processes. For example, airlines have hundreds, major ones even several thousands of take-offs every day. Thus, given this large number of similar repetitions, a single realization has a low impact, and the law of large numbers ensures that the average revenue per repetition is maximized and is also quite stable when using a risk-neutral model. As Barz (2007) points out, if the number of repetitions is too small, the assumption of both

costless insurance markets and perfect capital markets is necessary to convert the uncertain revenue stream into a certain one with the same present value.

Lancaster (2003) was the first to raise the issue of risk and revenue management. Interestingly, most authors cite consultants' experiences in practice to underline the relevance of risk-aversion. A consultant who worked with smaller airlines asked Weatherford (2004) about risk-averse capacity control. Barz (2007) reports that another consultant's clients were not comfortable with their risk-neutral revenue management system. They manually altered the forecast to obtain less aggressive (and risky) results. In an experiment conducted at a cruise line company, Singh (2011) observed that analysts' individual risk-aversion has a huge impact on their decisions when overriding a revenue management system's output. He attributes this behavior largely to their personality because they made decisions about exactly the same issues and possessed identical information. Levin et al. (2008) describe two other frequently cited examples. Event promoters may organize only a few large events per year in stadiums or concert halls that are very expensive to rent. Thus, their first priority would be to recover this fixed cost. In other industries, a manager's primary concern might be 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.

Weatherford (2004) was the first to bring revenue management and risk-aversion together. He modified the famous EMSR-b heuristic of Belobaba (1992) by substituting revenues with a risk-averse utility function. Barz and Waldmann (2007) also use an additive time-separable exponential utility function to model risk-aversion, but work with the dynamic programming formulation. They show that several well-known properties regarding the structure of an optimal policy carry over from the risk-neutral dynamic program. In her dissertation, Barz (2007) considers additional function types, such as atemporal utility functions that render the model less tractable. Moreover, she gives a general introduction to modeling risk attitudes with utility functions and how to integrate them into sequential decision problems. Feng and Xiao (2008) focus on an atemporal utility function

and show, among others, that higher risk-aversion leads to a more conservative policy. More specifically, a larger capacity is made available for cheaper products instead of waiting for customers who are willing to buy the expensive ones. Zhuang and Li (2011) also use an atemporal utility function but restrict themselves to only two products, which are sold one after another. In this context, they show structural properties regarding the degree of risk-aversion and demand volatility.

Apparently without being aware of Feng and Xiao's result, Huang and Chang (2011) constructed a straightforward heuristic that was directly derived from the exact dynamic program by artificially decreasing capacity's future value and, thus, selling more cheap products. Two variants of this approach are shown to yield more stable revenues than the standard risk-neutral dynamic program in the sense of a lower standard deviation. König and Meissner (2009a) have further refined this approach.

While the aforementioned research aims at optimizing utility, dispersion parameters, like standard deviation, are also widely used to evaluate the resulting policies. More recently, risk measures, such as Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR), have also been considered for evaluation in König and Meissner (2009a) as well as König and Meissner (2010). König and Meissner (2009b) remember the event promoter of Levin et al. (2008), who has to recover his fixed costs, and modify the dynamic program to minimize the Target Percentile Risk (TPR), which is the probability of obtaining less than a given target revenue. König and Meissner (2010) point out the need to "compute an optimal policy for the common risk measures, such as standard deviation, Value-at-Risk, or Conditional Value-at-Risk" and focus on the VaR. They determined the VaR-optimal policy by performing a binary search that repeatedly calculates TPR-optimal policies until the probability of failing the target revenue closely resembles the VaR's level.

VaR has its roots in finance and was first proposed by the global financial services firm J.P. Morgan Chase as an acceptability measure for a financial position with

random returns (Wozabal 2012). For a given probability level α , the VaR is the lowest revenue that will not be exceeded with a probability greater than or equal to α . Not least because of this simple definition as a quantile function, VaR is widely used in academia as well as in practice, where it has been incorporated into the Basel II and Solvency II regulations. However, despite its popularity, VaR has certain drawbacks. Obviously, it completely neglects the distribution of revenue below the quantile. Moreover, VaR does not belong to the important class of coherent risk measures, proposed in a seminal paper by Artzner et al. (1999), since it lacks certain desirable and intuitive properties. For example, it fails to satisfy the subadditivity property required for coherent risk measures, and, furthermore, the set of acceptable revenues may not be convex. Therefore, VaR should not be used to control multiple risky positions, as the individual control of each position does not allow for the control of their sum. The risk of a combined position may be deemed greater than the sum of the individual risks, and diversification may be penalized. In addition, VaR also fails to recognize the concentration of risks (Artzner et al. 1999). Owing to these shortcomings, VaR is increasingly being replaced by the Conditional Value-at-Risk (CVaR, also called Average VaR or Tail VaR). For continuous distributions, CVaR is simply the expectation below the VaR and, thus, considers “how bad is bad” (Artzner et al. 1999). In addition, it is a coherent risk measure (Pflug 2001). Despite these advantages, to the best of our knowledge, no work has been conducted that directly aims at optimizing CVaR in revenue management until now.

4 Modeling CVaR in Revenue Management

In this section, we present the new heuristic approach to optimizing CVaR in the context of revenue management. To this end, we state formal definitions of CVaR (Section 4.1) and discuss challenges when maximizing CVaR in a multistage revenue management environment together with a solution approach recently proposed in the literature (Section 4.2). Based on this, we develop a heuristic that is computationally efficient. To this end, we first reformulate the value function to

reduce its state space (Section 4.3). Furthermore, in Section 4.4, we show how the exploitation of the value function's structural properties facilitates solving the minimization problem, which occurs in each state. By reformulating this optimization problem as a continuous knapsack problem, we are able to give an efficient algorithm for the computation of the new approach.

4.1 General representations of CVaR

The first building block we need is a formal definition of CVaR. As the most intuitive definition of CVaR relies on the VaR, we start with a formal definition of the latter. Given a confidence level $\alpha \in [0; 1]$ and random variable Y on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with distribution function $F(y) = \mathbb{P}(Y \leq y)$, we use the following definition of VaR:

$$\text{VaR}_\alpha(Y) = \inf\{y: F(y) \geq \alpha\} = F^{-1}(\alpha), \quad (3)$$

where Y denotes a profit; that is, bigger values of Y are preferred. With (3), CVaR is now straightforwardly defined as

$$\text{CVaR}_\alpha(Y) = \mathbb{E}[Y: Y \leq \text{VaR}_\alpha(Y)] = \mathbb{E}[Y: Y \leq F^{-1}(\alpha)]. \quad (4)$$

From a theoretical point of view, (4) is valid only if considering probability spaces without atoms. But clearly, in the problem defined above, revenue is a discrete random variable for each given policy. Thus, we turn to the more common, but less intuitive definition

$$\text{CVaR}_\alpha(Y) = \inf\left\{\mathbb{E}[YZ]: \mathbb{E}[Z] = 1, 0 \leq Z \leq \frac{1}{\alpha}\right\} \quad (5)$$

with $\text{CVaR}_0 = \text{ess inf } Y$.

As Pflug and Pichler (2012) explain, the infimum in (5) is among all nonnegative random variables $Z \geq 0$ with expectation $\mathbb{E}[Z] = 1$ (densities), which satisfy the additional truncation constraint $Z \leq \frac{1}{\alpha}$. This representation is also known as CVaR's dual representation. To give some intuition on this formula, the Z can be viewed as weights that indicate whether an event (atom) falls below VaR_α and is thus included in or excluded from the CVaR's expectation. Loosely speaking,

what distinguishes (5) from (4), is that it is now possible to divide an atom and include only part of it in CVaR's expectation.

4.2 CVaR in a multistage environment

Given (4) and (5), CVaR_α is well defined for a given policy as described in Section 2. The authors mentioned in Section 3 use this in order to – mostly empirically – evaluate their policies. But how can we find a CVaR_α -optimal policy? At first, it might seem intuitive to simply substitute the expectation in (1) with (5). Although such a calculation of CVaR_α at every time step seems straightforward, it is unfortunately not the same as calculating CVaR_α over the whole selling horizon (see Artzner et al. 2007). The same holds for the optimization of CVaR_α in a multi-stage decision problem, as shown in this simple example: Assume that the level α is smaller than the probability that no customer arrives ($p_0(t)$) for every micro period. Then, obviously, only the case of no arrival in every period would be considered in such a step-wise CVaR, and, thus, the optimization. This completely ignores that the probability of no arrival throughout the whole selling horizon may be much smaller than α . Moreover, the number of micro periods is somewhat arbitrary, rendering a CVaR that is related to single, artificial micro periods less desirable.

As a recent result by Pflug and Pichler (2012) shows, the level α indeed cannot remain constant in multistage evaluations of CVaR, and changes with the probability of obtaining an atom included in the CVaR. One can use this result to obtain a value function that represents the maximum attainable CVaR_α from micro period t onwards:

$$\begin{aligned} V_t(X^{1:t-1}, Y^{1:t-1}, \alpha) &= \text{ess sup}_{x_t} \text{ess inf}_{Z_t} \sum_{i=0}^n p_i(t) \cdot Z_{ti} \cdot V_{t+1}(X^{1:t}, Y^{1:t}, \alpha \cdot Z_{ti}) \quad (6) \\ &\text{with } V_T(X^{1:T}, Y^{1:T}, \alpha) = r(X^{1:T}, Y^{1:T}), \end{aligned}$$

where the ess inf is over all Z with $\mathbb{E}[Z] = 1, 0 \leq Z \leq 1/\alpha$ and V_t depends on $X^{1:t} = (x_1, \dots, x_t)$, the complete history of past decisions and $Y^{1:t} = (y_1, \dots, y_t)$,

the history of past customer arrivals as well as $Z_t = (Z_{t0}, Z_{t1}, \dots, Z_{tn})^T$. The total revenue depends on the complete history of past decisions and customer arrivals and is denoted by $r(X^{1:T}, Y^{1:T})$. This shows how the stage-wise optimization can be retained while optimizing CVaR, but at the cost of a level α that is no longer constant, but becomes a random variable. The value of α in micro period $t > 1$ depends on $Y^{1:t-1}$, which is a random variable in earlier periods but becomes known in period t .

However, this formulation requires convexity of the action space $\mathcal{X}_t \ni x_t$. As capacity control's decisions on the acceptance of requests are discrete, and therefore $\mathcal{X}_t = \{0,1\}^n$, this condition is obviously not met and the aforementioned value function only serves as an approximation to the actual CVaR.

4.3 Reformulation of the value function

In the following section, we use Equation (6) as a starting point to construct our heuristic. Calculating (6) in the context of revenue management poses two challenges: the intractably large state space and the computationally expensive minimization problem for each stage.

The first challenge, the high dimensionality of its state space, results from the inclusion of the decision history $X^{1:T}$ as well as the past demand realizations $Y^{1:T}$, which are necessary to calculate $r(X^{1:T}, Y^{1:T})$ at the final stage. Fortunately, the state space can be simplified as follows: First, observe that in (6) past decisions and demand realizations are only needed to calculate the revenue obtained in the past and to determine what additional requests can be accepted; for the latter, knowing the remaining capacity is sufficient. Moreover, revenues can be considered directly at each stage owing to CVaR's translation equivariance, namely:

$$\text{CVaR}_\alpha(Y + a) = \text{CVaR}_\alpha(Y) + a,$$

for a random variable Y and $a \in \mathbb{R}$. Thus, similar to (1), it suffices to include the remaining capacity in the state space in addition to the level α and (6) can be written as

$$V_t(c, \alpha) = \text{ess sup}_{x_t} \text{ess inf}_{Z_t} \sum_{i=1}^n p_i(t) \cdot Z_{ti} \cdot (x_{ti} \cdot r_i + V_{t+1}(c - x_{ti}, \alpha \cdot Z_{ti})) + p_0(t) \cdot Z_{t0} \cdot V_{t+1}(c, \alpha \cdot Z_{t0}), \quad (6)$$

subject to the boundary conditions: $V_t(c, \alpha) = -\infty$ for $c < 0$ and $V_T(c, \alpha) = 0$ for $c \geq 0$. As in (2), the firm should accept a request for product i if and only if

$$r_i \geq V_{t+1}(c, \alpha) - V_{t+1}(c - 1, \alpha),$$

where, as stated above, the value of α becomes known in period t .

4.4 Efficient implementation

For a straightforward implementation of (6), we employed the *fmincon* function, which is part of the Optimization Toolbox for MATLAB, in order to solve the optimization problem in each state. The function performs a constrained minimization without requiring any information about the objective function and therefore takes no advantage of any inherent structural properties.

To calculate (6) more efficiently, we exploit structural properties that come to light when we transform the value function $V_t(c, \alpha)$, which seems to possess no clear structure besides the obvious monotonicity. In contrast, the function

$$\tilde{V}_t(c, \alpha) := \alpha \cdot V_t(c, \alpha) \quad (7)$$

is also convex in α (see Pflug and Pichler 2012). Additionally, in the context of capacity control, $\tilde{V}_t(c, \alpha)$ is piecewise linear in α due to the discrete distribution of revenues. This structure can be used to considerably simplify the optimization problem and, as it is necessary to solve it in each state, this has a large impact on the overall run-time. Below, we present an algorithm that efficiently solves the minimization problem to optimality.

First, by substituting (6) into (7), we obtain

$$\tilde{V}_t(c, \alpha) = \text{ess sup}_{x_t} \text{ess inf}_{Z_t} \sum_{i=1}^n p_i(t) (\alpha \cdot Z_{ti} \cdot x_{ti} \cdot r_i + \tilde{V}_{t+1}(c - x_{ti}, \alpha \cdot Z_{ti})) + p_0(t) \cdot \tilde{V}_{t+1}(c, \alpha \cdot Z_{t0}).$$

Now, let

$$f_t(c, \alpha, x_t, Z_t) := \sum_{i=1}^n p_i(t) \cdot (\alpha \cdot Z_{ti} \cdot x_{ti} \cdot r_i + \tilde{V}_{t+1}(c - x_{ti}, \alpha \cdot Z_{ti})) + p_0(t) \cdot \tilde{V}_{t+1}(c, \alpha \cdot Z_{t0}). \quad (8)$$

Moreover, as we consider only discrete random variables with compact support, the ess inf is in fact a minimum. Therefore, at each state (c, α) and for each decision x_t , we have to solve the minimization problem

$$\min_{Z_t} f_t(c, \alpha, x_t, Z_t) \quad (9)$$

$$\text{s. t.} \quad p_t^T Z_t = 1 \quad (10)$$

$$0 \leq Z_{ti} \leq \frac{1}{\alpha}, i = 0, \dots, n \quad (11)$$

$$Z_t \in \mathbb{R}^{n+1} \quad (12)$$

where $p_t = (p_0(t), \dots, p_n(t))^T$. With $\tilde{Z}_t := \alpha \cdot Z_t$, we combine the level α and Z_t to rewrite (8) as $\tilde{f}_t(c, x_t, \tilde{Z}_t) := \sum_{i=1}^n p_i(t) \cdot (\tilde{Z}_{ti} \cdot x_{ti} \cdot r_i + \tilde{V}_{t+1}(c - x_{ti}, \tilde{Z}_{ti})) + p_0(t) \cdot \tilde{V}_{t+1}(c, \tilde{Z}_{t0})$. Using this, the abovementioned problem can be rewritten as

$$\min_{\tilde{Z}_t} \tilde{f}_t(c, x_t, \tilde{Z}_t) \quad (13)$$

$$\text{s. t.} \quad p_t^T \tilde{Z}_t = \alpha \quad (14)$$

$$0 \leq \tilde{Z}_{ti} \leq 1, i = 0, \dots, n \quad (15)$$

$$\tilde{Z}_t \in \mathbb{R}^{n+1}, \quad (16)$$

where \tilde{f}_t is piecewise linear in $\tilde{Z}_{ti}, i = 0, \dots, n$, which is obvious when \tilde{f}_t is written as $\tilde{f}_t(c, x_t, \tilde{Z}_t) := \sum_{i=1}^n p_i(t) \cdot \tilde{f}_{ti}(c, x_t, \tilde{Z}_{ti}) + p_0(t) \cdot \tilde{f}_{t0}(c, x_t, \tilde{Z}_{t0})$ with $\tilde{f}_{ti}(c, x_t, \tilde{Z}_{ti}) := \tilde{Z}_{ti} \cdot x_{ti} \cdot r_i + \tilde{V}_{t+1}(c - x_{ti}, \tilde{Z}_{ti}), i = 1, \dots, n$ and $\tilde{f}_{t0}(c, x_t, \tilde{Z}_{t0}) := \tilde{V}_{t+1}(c, \tilde{Z}_{t0})$ due to the piecewise linearity of \tilde{V}_{t+1} .

As $x_t \in \mathcal{X}_t = \{0,1\}^n$, the evaluation of $\tilde{f}_{ti}(c, x_t, \tilde{Z}_t)$ only involves $\tilde{V}_{t+1}(c - 1, \cdot)$ if requests for product i are accepted ($x_{ti} = 1$) and $\tilde{V}_{t+1}(c, \cdot)$ if requests for product i are rejected ($x_{ti} = 0$) in period t , including the case of no customer arrival. Subsequently, we denote by $k_{a1} = 0 < k_{a2} < \dots < k_{an_a} = 1$ the set of points, where $\tilde{V}_{t+1}(c - 1, \cdot)$ is not differentiable, including the endpoints 0 and 1.

Likewise, we denote by $k_{r1} = 0 < k_{r2} < \dots < k_{rn_r} = 1$ the set of points, where $\tilde{V}_{t+1}(c, \cdot)$ is not differentiable, again including the endpoints 0 and 1. Thus, with s_{ij} , we can denote the slope of \tilde{f}_{ti} in segment j that is between $k_{a,j}$ and $k_{a,j+1}$ for requests accepted ($x_{ti} = 1$) and between $k_{r,j}$ and $k_{r,j+1}$ for requests rejected ($x_{ti} = 0$). This structure enables us to partition each variable \tilde{Z}_{ti} into variables \tilde{Z}_{tij} corresponding to segment j of \tilde{f}_{ti} and write \tilde{f}_{ti} as $\tilde{f}_{ti}(c, x_t, \tilde{Z}_{ti}) = \sum_{j=1}^{n_a} \tilde{Z}_{tij} \cdot s_{ij}$ with $\sum_{j=1}^{n_a} \tilde{Z}_{tij} = \tilde{Z}_{ti}$, for $x_{ti} = 1$, where n_a is replaced by n_r for $x_{ti} = 0$. From the convexity of \tilde{V}_{t+1} , it follows that $s_{ij} < s_{i,j+1}$. Thus, (9)–(12) can be expressed as the following linear program:

$$\min_{\tilde{Z}_t} \sum_{i:x_{ti}=1} \sum_{j=1}^{n_a} \tilde{Z}_{tij} \cdot (p_i(t) \cdot s_{ij}) + \sum_{i:x_{ti}=0} \sum_{j=1}^{n_r} \tilde{Z}_{tij} \cdot (p_i(t) \cdot s_{ij}) \quad (17)$$

$$\text{s. t.} \quad 0 \leq \tilde{Z}_{tij} \leq k_{a,j+1} - k_{a,j} \quad \forall i, j: x_{ti} = 1 \quad (18)$$

$$0 \leq \tilde{Z}_{tij} \leq k_{r,j+1} - k_{r,j} \quad \forall i, j: x_{ti} = 0 \quad (19)$$

$$\sum_{i:x_{ti}=1} \sum_{j=1}^{n_a} \tilde{Z}_{tij} \cdot p_i(t) + \sum_{i:x_{ti}=0} \sum_{j=1}^{n_r} \tilde{Z}_{tij} \cdot p_i(t) = \alpha \quad (20)$$

Moreover, Problem (17)–(20) is in fact a continuous knapsack problem where \tilde{Z}_{tij} is the number of items selected of type (t, i, j) and $p_i(t)$ represents the items' weight. Constraints (18) and (19) represent the amount available of each item, and constraint (20) is the weight restriction of the knapsack. Thus, Problem (17)–(20) can be solved efficiently by using a simple greedy procedure (Dantzig 1957).

The basic idea of such a procedure is to sort the items according to their relative value per unit weight, which is $\frac{p_i(t) \cdot s_{ij}}{p_i(t)} = s_{ij}$ in this case. As we have a minimization problem, the items are subsequently considered in the order of increasing relative value. As much as possible is selected from each item with respect to constraints (18) and (19). The procedure terminates when our knapsack is full and condition (20) is satisfied.

We provide the formal algorithm, which we call *lowest-ascent*, in the following. Steps 1–3 involve some initializations. The core of the algorithm is Step 4. Here, we iterate over the slopes until the knapsack is full; that is, condition (20) is

satisfied with $\mathbb{E}[\tilde{Z}_t] = \alpha$. In each iteration k , we select as much as possible from every item with relative value $s_{(k)}$ (Steps 4.1 and 4.2). This amount is restricted by the amount available (Steps 4.2.1 and 4.2.2) and, for the last item, by the knapsack's maximum weight (Step 4.2.3), whichever is less. As we are only interested in the \tilde{Z}_{ti} to obtain a solution for Problem (13)–(16), we directly add the increase to \tilde{Z}_{ti} (Step 4.2.4).

Algorithm *lowest-ascent*:

1. $\tilde{Z}_{ti} = 0, i = 0, \dots, n$
2. $k = 1$
3. *sort all unique slopes s_{ij} in ascending order $s_{(1)} < s_{(2)} < \dots < s_{(N)}$*
4. *while $\mathbb{E}[\tilde{Z}_t] < \alpha$*
 - 4.1. $S = \{(i, j) : s_{ij} = s_{(k)}\}$
 - 4.2. *for $(i, j) \in S$*
 - 4.2.1. *if $x_{ti} = 0$: $\Delta_1 = k_{r,j+1} - k_{rj}$*
 - 4.2.2. *else: $\Delta_1 = k_{a,j+1} - k_{aj}$*
 - 4.2.3. $\Delta_2 = \frac{\alpha - \mathbb{E}[\tilde{Z}_t]}{p_i(t)}$
 - 4.2.4. $\tilde{Z}_{ti} = \tilde{Z}_{ti} + \min(\Delta_1, \Delta_2)$
 - 4.2.5. *if $\mathbb{E}[\tilde{Z}_t] = \alpha$: break*
 - 4.3. $k = k + 1$

5 Simulation experiment design

In this section, we describe the design of simulation experiments used to compare the new approach developed in Section 4 to standard benchmark approaches. All implementations were done in MATLAB version R2012b and were run on a PC with a 2.8 GHz Intel Core i7 processor and 8 GB of RAM, running on Microsoft Windows Server 2008 Enterprise SP2 64 bit. We outline the approaches and control mechanisms considered in the simulation study in Section 5.1 and describe the settings considered in Section 5.2. In Section 5.3, we determine the grid size for α based on an analysis of how different grid sizes influence the accuracy and

runtime of the algorithm. The subsequent Section 6 contains the results of the simulation study.

5.1 Control mechanisms

In the simulation study, we compare the new approach to several standard control mechanisms. The following mechanisms are considered:

- *CVAR* denotes the new control mechanism developed in Section 4. We discretized the state space and used an equidistant grid with a grid size of 5%, i.e. $\alpha \in \{0, 0.05, 0.1, \dots, 1\}$, together with a linear interpolation in order to compute the value function via dynamic programming. In Section 5.3, we show that this is a reasonable compromise between accuracy and runtime.
- *EV* implements the control mechanism given by Equation (2) and uses the opportunity cost calculated by the classical dynamic program (Equation (1)). This mechanism is optimal in that it maximizes the revenue's expected value.
- *FCFS* is a very simple mechanism that accepts all requests in a first come, first served manner until no more capacity is left.
- *EXPOST* is a mechanism that uses perfect hindsight information on incoming demand to pick ex post the requests to accept from each demand stream. Thus, it calculates the maximum revenue that can be obtained from each demand stream and thereby maximizes expected revenue and Conditional Value-at-Risk. The obtained values serve as an upper bound for all other mechanisms.

5.2 Settings

The design of our simulation experiments is based on a classical example by Lee and Hersh (1993). Their setting is widely used in revenue management (see, e.g., Subramanian et al. 1999), especially in the context of risk-aversion (see, e.g., Barz

2007, Barz and Waldmann 2007 and König and Meissner 2009a). The example consists of an airplane with a capacity of $C = 10$ seats and four products (booking classes) $i \in I = \{1, \dots, 4\}$ with revenues of $r_1 = 200, r_2 = 150, r_3 = 120$ and $r_4 = 80$. As it is often the case in practice, demand for more expensive products tends to arrive later. The selling horizon consists of 30 micro periods and is partitioned into five intervals. Our first setting is the original setting from Lee and Hersh (1993) and, thus, denoted as *Setting Lee/Hersh*. The request probabilities for each interval are shown in Table 1, where $p_i(t)$ denotes the probability of a request for product i in micro period t . These probabilities reflect the fact that in passenger air transportation, higher value demand (often business customers) tends to arrive later in the selling horizon.

Table 1 Request probabilities for *Setting Lee/Hersh*

Request probability	Micro period t				
	1 – 5	6 – 12	13 – 19	20 – 26	27 – 30
$p_1(t)$	0.08	0.06	0.10	0.14	0.15
$p_2(t)$	0.08	0.06	0.10	0.14	0.15
$p_3(t)$	0.14	0.14	0.10	0.16	0
$p_4(t)$	0.14	0.14	0.10	0.16	0

To further evaluate the robustness of our results, we use two additional variants of this setting that differ in their arrival of demand over time. In the first variant, denoted as *Setting LBH*, demand follows the classical low-before-high assumption and customers demanding lower value products arrive strictly before customers asking for more expensive products. Table 2 shows the request probabilities. We chose the probabilities such that total expected demand for each product is the same as in Lee and Hersh’s example. The length of the time intervals is set such that the total number of micro periods remains at 30 and the demand intensity increases towards the end of the selling horizon. Again, this effect is widely observed in airline practice.

Table 2 Request probabilities for *Setting LBH*

Request probability	Micro period t			
	1 – 10	11 – 20	21 – 25	26 – 30
$p_1(t)$	0	0	0	0.62
$p_2(t)$	0	0	0.62	0
$p_3(t)$	0	0.35	0	0
$p_4(t)$	0.35	0	0	0

In the third setting, we consider the time-homogeneous arrival probabilities shown in Table 3 (*Setting Flat*). Again, 30 micro periods are used and the probabilities are set such that total expected demand is the same as in *Setting Lee/Hersh*.

Table 3 Request probabilities for *Setting Flat*

Request probability	Micro period t
	1 – 30
$p_1(t)$	0.1033
$p_2(t)$	0.1033
$p_3(t)$	0.1167
$p_4(t)$	0.1167

To evaluate our approach's performance, we generated 10,000 customer streams for each setting and used the abovementioned mechanisms to decide on the acceptance of requests and calculate the resulting revenue for each stream.

5.3 Grid size and runtime

We started the numerical experiments with a preliminary investigation of the grid size used in *CVAR* to discretize the state space with regard to α . For this purpose, we performed *CVAR* with five different grid sizes, 0.01, 0.05, 0.10, 0.15, and 0.20, to compute the value function given by Equation (6). For each grid size, we evaluated the demand streams 101 times, optimizing the CVaR at levels $\alpha \in$

$\{0,01, 0.02, \dots, 1\}$. Figure 1 shows the CVaRs obtained in *Setting Lee/Hersh*. As the results for other settings are very similar, they were omitted.

Whereas very low and high values of α have similar results, they vary considerably for intermediate levels. For values of α between about 0.1 and 0.5, the difference in CVaR is up to nearly 5%. Obviously, a finer grid size leads to higher CVaRs, but the difference between a grid size of 0.05 and 0.01 is negligible for all values of α .

In addition to these results, which were obtained using our *lowest-ascent* algorithm from Section 4.4, we also tested a straightforward implementation using the *fmincon* function provided by MATLAB's Optimization Toolbox version 6.2.1 with a grid size of 0.05. Compared to our algorithm with the same grid size, this implementation performed slightly worse in respect to the attained CVaR and was therefore excluded from further analysis. The weaker performance is probably due to the numerical approximations used by *fmincon*.




Figure 1 Attained CVaR (*Setting Lee/Hersh*)

Lowest-ascent's advantage over *fmincon* is even more obvious when runtimes are considered (Table 4). For a grid size of 0.05, *lowest-ascent* is about 90x faster

than *fmincon*. Jointly considering the attained CVaR and runtime, we think that a grid size of 0.05 offers a reasonable tradeoff and continue to use it in the following. Nonetheless, the new approach requires the solution of a continuous knapsack problem in every state. This leads to longer runtimes when compared to the risk-neutral dynamic program. As the settings only differ in the arrival probabilities, and therefore the computational complexity of each approach (risk-averse and risk-neutral) is the same in all settings, we explicitly state the runtimes only for *Setting Lee/Hersh*.

Please note that the values given in Table 4 refer to the time necessary to calculate the value function, which is necessary only once and can be done upfront, prior to the selling horizon. In all the approaches, the same time is needed to actually decide on whether or not to accept an arriving request. Moreover, to obtain results that are easy to compare and owing to very short runtimes, our implementation does not use parallelization. However, due to the specific structure of the dynamic program, the algorithm could easily be implemented using parallel computing. Consequently, the number of threads could easily reach the full initial capacity.

Table 4 Runtimes in *Setting Lee/Hersh*

Control mechanism	Grid size	Runtime [mm:ss.ms]
<i>CVAR</i>	0.01	10:17.514
<i>CVAR</i>	0.05	00:24.882
<i>CVAR</i>	0.10	00:06.490
<i>CVAR</i>	0.15	00:03.260
<i>CVAR</i>	0.20	00:02.000
<i>CVAR with fmincon</i>	0.05	35:51.545
<i>EV</i>	–	00:00.156

6 Numerical Results

In this section, we present the main results of an extensive simulation study that compares the new approach developed in Section 4 to standard benchmark approaches. In Section 6.1, we compare the CVaR obtained by the new approach with the benchmark mechanisms' results. In Section 6.2, we discuss the influence of demand's temporal distribution and compare the mechanisms to results obtained with perfect hindsight information. Finally, we take a look at the capacity utilization (Section 6.3) and the tradeoff between risk-aversion (i.e. maximizing CVaR) and expected revenue (Section 6.4).

6.1 Comparison of control mechanisms' CVaRs

In this section, we compare the revenue performance of the new approach with the benchmarks in each of the three settings described above. We first compare the CVaR of the revenues obtained with our approach with the benchmarks' CVaR. Second, focusing on the influence of demand's arrival order, we use the same method to compare the revenues obtained in the different settings. Finally, we take a look at capacity utilization to investigate how the observed revenues are attained.

Figures 2–4 show the CVaR obtained using the four control mechanisms outlined in Section 5.1. For each setting, we calculated the CVaR_α for $\alpha \in \{0, 0.05, 0.1, \dots, 1\}$. As the acceptance decisions of *EV*, *FCFS*, and *EXPOST* do not depend on the level α , we first processed all demand streams using these mechanisms and subsequently calculated the CVaR_α using the resulting per-stream revenues. In contrast, *CVAR*'s decisions depend on the level α and we have to process the streams for each α to determine the distribution of total revenue and calculate the CVaR.

CVAR clearly outperforms *FCFS* for all levels of α in all three settings and outperforms *EV* for $\alpha \leq 0.5$ in *Setting Lee/Hersh* (Figure 2) and *Setting Flat* (Figure 4). In *Setting LBH* (Figure 3), *CVAR* is superior to *EV* for all $\alpha \leq 0.7$. For higher values of α , there is essentially no difference. This behavior is intuitive,

because as $\alpha \rightarrow 1$, CVaR_α approaches the expected value and CVAR 's policy increasingly resembles that of EV . For low levels of α (approximately below 0.2), FCFS performs better than EV and almost as well as CVAR . This is due to the fact that a low confidence level means considering only the worst outcomes, which usually are the ones with no or very weak demand. Thus, with decreasing α , CVAR accepts more requests upfront and becomes more similar to FCFS , which accepts all requests. In contrast, EV denies some requests for products with lower revenues and rather waits for requests for products with higher revenues. Therefore, FCFS can be considered a policy for an extremely risk-averse decision maker. Please note that, while CVAR is formally equivalent to EV for $\alpha = 1$, it is not fully equivalent to FCFS for $\alpha = 0$ because CVAR involves waiting for high revenue requests that will arrive for sure. In sum, in realistic settings, FCFS and EV can be considered two extremes between one can vary by choosing a level $\alpha \in [0,1]$ for CVAR .




Figure 2 Attained CVaR (*Setting Lee/Hersh*)




Figure 3 Attained CVaR (*Setting LBH*)




Figure 4 Attained CVaR (*Setting Flat*)

6.2 The influence of demand's temporal distribution

Next, we focus on the influence of demand's arrival order. For each control mechanism, we compare the CVaR obtained when demand is low-before-high (*Setting LBH*), somewhere in-between (*Setting Lee/Hersh*), and time-homogeneous (*Setting Flat*). Figures 5–7 show the resulting CVaR for *EV*, *FCFS*, and *CVAR* relative to the CVaR obtained with *EXPOST*. This benchmark is widely used in numerical studies in revenue management, even though it disposes of perfect hindsight demand information and, therefore, provides only an easy-to-calculate upper bound on the expected revenue. This bound is usually not attainable through control methods disposing of only distributional information. Thus, the gap between a control mechanism and *EXPOST* usually consists of two components: the value of information and what we call the optimality gap that stems from the mechanisms' heuristic nature. Moreover, in this context, *EXPOST* is useful because it does not depend on the setting; it ignores demand's temporal distribution and only uses information on total expected demand, which is equal in all three settings. As CVaR equals the expected value at $\alpha = 1$, the expected revenue is displayed at the right of all figures.

Figure 5 displays the performance of *EV*, which is optimal in terms of expected revenue. Thus, at $\alpha = 1$, we have the optimal expected revenue that can be obtained. Here, the difference between *EV* and *EXPOST* is solely the value of information and there is no optimality gap. Not surprisingly, revenues decline and this gap increases when lower value products are requested earlier and capacity control becomes more challenging. Regarding smaller values of α , the difference between *EV* and *EXPOST* increases because, besides the information value, there is also an optimality gap between them as we still follow a policy that optimizes the expected value instead of the evaluated CVaR depicted in the figure.

The results for *CVAR* are shown in Figure 6. *CVAR* performs best in the time-homogeneous case (*Setting Flat*), where it delivers at least 96% of *EXPOST*. Its performance is slightly lower in *Setting Lee/Hersh*, but still quite impressive; it

delivers at least 95% of *EXPOST*. *CVAR* performs worst in *Setting LBH* where demand arrives strictly low-before-high, delivering at least 92% of *EXPOST*. When comparing *CVAR* and *EV*, it is important to remember that Figures 5 and 6 display the same values for the level $\alpha = 1$. However, whereas the relative CVaRs obtained through *EV* decline with decreasing α (from right to left), *CVAR* yields constant results. Only *Setting Flat* shows a slight decline of about 0.5% from $\alpha = 1$ to $\alpha = 0.9$. Beginning at about $\alpha = 0.3$, the relative CVaR rises until it reaches 100% in all three settings at $\alpha = 0$. Regarding the optimality of the policy calculated with *CVAR*, the comparison to *EXPOST* is only of limited explanatory power because, with our heuristic approach, we are not able to determine an optimal policy or confirm that our policy is optimal. Therefore, we do not know how much of the difference is due to the information value. However, for $\alpha = 1$, we know that the difference equals the information value. As the difference is the biggest for $\alpha = 1$, we are quite confident that the policy is quite good, although we do not know how the value of information declines with decreasing α .

For the sake of completeness, Figure 7 shows the results for *FCFS*. In general, the picture is quite similar to that of *CVAR*, but the attained CVaRs are considerably lower at about 90%, 87%, and 77% for the three settings.




Figure 5 Percentage of *EXPOST* (*EV*)




Figure 6 Percentage of *EXPOST* (*CVAR*)




Figure 7 Percentage of *EXPOST* (*FCFS*)

The value of $\alpha = 0$ is a special case. According to (5), the $CVaR_0$ only reflects the worst outcome, that is the simulation run with the least revenue. Thus, the result is more or less arbitrary, even with a very large number of simulation runs. In *Settings Lee/Hersh* and *LBH*, as expected, *EV* obtains quite a small revenue in the worst customer stream compared to *EXPOST*'s revenue in the stream where *EXPOST* obtains the lowest revenue (The worst streams are not necessarily the same for different control mechanisms). But surprisingly, *EXPOST* and *EV* obtain identical revenues in the worst stream of *Setting Flat*. A detailed analysis shows that these worst streams for *EXPOST* and *EV* are identical and only contain 3 customer requests. Obviously, *EXPOST* accepts all 3 requests, even though they are for low value products. But *EV* also accepts the requests, because they arrive quite late in the selling horizon. Thus, we get identical values for $CVaR_0$ for both methods, although this result is quite unusual because the probability that such a stream is generated is extraordinarily low.

6.3 Capacity utilization

To investigate how *CVAR* achieves its results, we consider the mean capacity utilization for each setting, shown in Figure 8. At first, we notice that the capacity utilization declines in α : The less risk-averse we become, that is, the higher α , the

longer *CVAR* waits for higher revenue requests that may (or may not) arrive in the future. More precisely, early arriving lower revenue requests are rejected and the lowest capacity utilization is attained for $\alpha = 1$, which is the same policy as *EV*. In *Setting Flat* and *Setting Lee/Hersh*, the capacity utilization declines rather smoothly in α , whereas in *Setting LBH*, it seems to decline in three steps. This seemingly odd behavior is caused by demand arriving strictly low before high. The average number of accepted requests in all three settings is shown in Figure 9–Figure 11 and Table 5 in the Appendix. Figure 10 reveals that in *Setting LBH*, *CVAR* seems to follow some kind of protection level approach, at least regarding the earliest (lowest revenue) product 4. The steps correspond to (from left to right) accepting a maximum of 3, 2, 1 or 0 requests for that product. The effect of this protection level approach for product 4 on total capacity utilization is larger than the influence of the other products, and determines the prevalent structure of capacity utilization. Note that average numbers are smaller than these integers as they reflect the probability that the respective number of requests will actually arrive.




Figure 8 Capacity Utilization




Figure 9 Average number of accepted requests (CVAR, Setting Lee/Hersh)

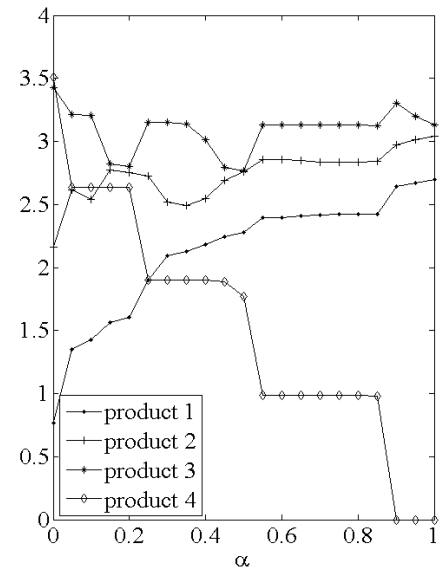


Figure 10 Average number of accepted requests (CVAR, Setting LBH)




Figure 11 Average number of accepted requests (CVAR, Setting Flat)

6.4 Tradeoff between risk-aversion and expected revenue

Although our primary goal is to maximize CVaR for a certain level of α , we briefly consider revenue's standard deviation in this section. Figures 12–14 show scatter plots of the standard deviation against the mean of total revenues for the three settings considered. Each of the 21 points in a plot represents employing

CVAR with a specific value for α and computing the sample mean and standard deviation of total revenues across all customer streams. In addition, the values obtained with *EV* and *FCFS* are indicated. As already discussed in Section 6, they are identical to the values obtained with *CVAR* for $\alpha = 1$ and $\alpha = 0$, respectively. We notice a large reduction in the standard deviation when using *CVAR* with $\alpha < 1$ compared to *EV* ($\alpha = 1$). As the confidence level α declines, the standard deviation also declines further. However, this trend only holds for $\alpha \geq 0.2$. For smaller values of α , the standard deviation increases again and reaches the value of *FCFS* for $\alpha = 0$. While capacity control decisions for intermediate values of α apparently smooth the revenues obtained from the stochastic customer streams, this does not hold for very low values of α . Here, the worst outcomes are increasingly considered and capacity control increasingly resembles *FCFS* which accepts all requests until the capacity is fully occupied. Using *FCFS*, we obtain simply the mean and the standard deviation of the sum of the first ten requests' prices.

Not surprisingly, all three figures clearly show that the optimization of CVaR for a level $\alpha \neq 1$ and, thus, the reduction in standard deviation comes at the price of a decrease in expected revenue. Note that while the points are quite evenly spread in *Setting Lee/Hersh* and *Setting Flat*, they are clustered into about four groups in *Setting LBH*. This again reflects that there are only a few major policy changes as α is varied.




Figure 12 Comparison of expected value and standard deviation (*Setting Lee/Hersh*)




Figure 13 Comparison of expected value and standard deviation (*Setting LBH*)




Figure 14 Comparison of expected value and standard deviation (*Setting Flat*)

7 Conclusions

As revenue management increasingly spreads to industries other than the large airlines where it was born some 30 years ago, new challenges arise. One issue is the assumption of risk-neutrality that lies at the heart of traditional revenue management. While this assumption is, on average, perfectly appropriate in the long run, based on a high number of flight departures, hotel room bookings, etc., it does not take into account that decision makers may prefer stable results in the

short term rather than accepting the volatility caused by risk-neutral models. While users' acceptance of automated systems is always an issue in practice, we observed in our consulting practice that even highly trained analysts with a mathematical background (rightly?) mistrust risk-neutral systems. However, as Singh (2011) observed in an experiment conducted at a cruise line company, analysts following only their intuition make very different and contradicting decisions when overriding a revenue management system's output.

As a result, risk-averse models have been considered in revenue management since about a decade ago and a variety of risk indicators have been used. However, although research on risk and revenue management has recently gained momentum, our literature review shows that the body of scientific publications is still quite limited.

This paper contributes to that research by proposing – to the best of our knowledge – the first approach that directly aims at maximizing CVaR, an intuitively appealing risk measure that is widely used in finance due to its desirable theoretical properties. The exploitation of the underlying value function's structure allows us to reformulate the inherent optimization problem as a continuous knapsack problem and provide an algorithm that allows for an efficient implementation. Thus, the runtime is reduced to a fraction of the time needed by a straightforward implementation that uses standard optimization functions. In a simulation study based on an example widely used in the literature, we show that the new approach delivers very good results. Moreover, we give insights into how these results are obtained and discuss the approach's relationship to widely used benchmark approaches.

In our opinion, there are two main avenues for further research on this topic. The first specifically relates to our approach and includes extensions to other risk measures that admit a Choquet representation, and, thus, are closely related to CVaR. Moreover, it may be worthwhile to obtain upper bounds that are tighter than our variant of the standard benchmark using hindsight information. The second avenue relates to research on risk-averse revenue management in general,

where the consideration of a single resource and so-called independent demand is standard. This is justified by risk-aversion being most relevant to smaller firms. Moreover, these assumptions allow for focusing on the effects of risk-aversion in models with an improved mathematical tractability. Nonetheless, resource networks and customer choice are highly relevant in some industries and increasingly considered in research on classical risk-neutral revenue management. Therefore, we think that bringing these elements and risk-aversion together would be an interesting avenue for future research. The extension of our approach to resource networks is quite straightforward, although all the difficulties known from classical network revenue management will come into play. Similarly, the inclusion of customer choice is a possibility but the challenge will be to develop a model that is computationally tractable.

Appendix

Table 5 CVAR's average number of accepted requests per product

Setting	Product	α																				
		0	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4	0.45	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9	0.95	1
<i>Lee/Hersh</i>	1	2.04	2.07	2.12	2.17	2.24	2.29	2.34	2.39	2.44	2.48	2.53	2.57	2.61	2.66	2.69	2.72	2.75	2.78	2.8	2.83	2.84
	2	2.01	2.04	2.09	2.14	2.2	2.25	2.29	2.34	2.39	2.43	2.47	2.51	2.54	2.58	2.61	2.64	2.66	2.68	2.7	2.71	2.72
	3	2.89	2.91	2.93	2.96	2.98	3	3.03	3.04	3.06	3.07	3.08	3.08	3.11	3.11	3.11	3.1	3.08	3.06	3.04	3.02	3
	4	2.91	2.83	2.7	2.55	2.39	2.24	2.1	1.97	1.84	1.72	1.6	1.49	1.35	1.23	1.15	1.07	1.01	0.97	0.93	0.9	0.88
<i>LBH</i>	1	0.77	1.35	1.43	1.56	1.6	1.9	2.09	2.13	2.18	2.24	2.28	2.4	2.4	2.41	2.42	2.42	2.42	2.42	2.64	2.67	2.7
	2	2.16	2.61	2.54	2.78	2.75	2.73	2.52	2.49	2.55	2.69	2.76	2.86	2.86	2.85	2.84	2.84	2.84	2.84	2.97	3.01	3.04
	3	3.43	3.21	3.2	2.82	2.8	3.15	3.15	3.14	3.01	2.8	2.77	3.13	3.13	3.13	3.13	3.13	3.13	3.12	3.3	3.2	3.13
	4	3.51	2.64	2.64	2.64	2.64	1.9	1.9	1.9	1.9	1.88	1.77	0.99	0.99	0.99	0.99	0.99	0.99	0.98	0	0	0
<i>Flat</i>	1	2.3	2.31	2.33	2.37	2.41	2.44	2.48	2.52	2.56	2.59	2.63	2.66	2.69	2.72	2.75	2.78	2.81	2.83	2.85	2.87	2.89
	2	2.33	2.34	2.37	2.4	2.44	2.47	2.51	2.54	2.58	2.61	2.64	2.68	2.7	2.74	2.77	2.79	2.81	2.83	2.84	2.86	2.87
	3	2.59	2.6	2.62	2.65	2.68	2.71	2.73	2.76	2.78	2.81	2.83	2.85	2.86	2.89	2.91	2.91	2.9	2.88	2.86	2.84	2.82
	4	2.63	2.6	2.52	2.42	2.29	2.19	2.08	1.96	1.85	1.73	1.63	1.53	1.44	1.31	1.21	1.13	1.07	1.03	1	0.98	0.96

References

- Artzner P, Delbaen F, Eber JM, Heath D (1999) Coherent measures of risk. *Math Finance*, 9(3):203–228
- Artzner P, Delbaen F, Eber JM, Heath D, Ku H (2007) Coherent multiperiod risk adjusted values and Bellman's principle. *Ann of Oper Res*, 152(1):5–22
- Barut M, Sridharan V (2005) Revenue management in order-driven production systems. *Decis Sci* 36(2):287–316
- Barz C (2007) Risk-averse capacity control in revenue management. Springer, Berlin
- Barz C, Waldmann KH (2007) Risk-sensitive capacity control in revenue management. *Math Methods of Oper Res*, 65(3):565–579
- Belobaba PP (1992) Optimal vs. heuristic methods for nested seat allocation. Proc of AGIFORS Reserv and Yield Manag Study Group. Brussels, pp 28–53
- Bharill R, Rangaraj N (2008) Revenue management in railway operations: A study of the Rajdhani Express, Indian Railways. *Transp Res Part A: Policy and Pract*, 42(9):1195–1207
- Carroll WJ, Grimes RC (1995) Evolutionary change in product management: Experiences in the car rental industry. *Interfaces* 25(5):84–104
- Ciancimino A, Inzerillo G, Lucidi S, Palagi L (1999) A mathematical programming approach for the solution of the railway yield management problem. *Transp Sci* 33(2):168–181
- Chiang WC, Chen JCH, Xu X (2007) An overview of research on revenue management – current issues and future research. *Int J of Revenue Manag* 1(1): 97–128
- Cross RG (1997) Revenue management: Hardcore tactics for market domination. Broadway Books, New York
- Dantzig GB (1957) Discrete-variable extremum problems. *Oper Res* 5(2):266–288
- Defregger F, Kuhn H (2007) Revenue management for a make-to-order company with limited inventory capacity. *OR Spectr*, 29(1):137–156
- Feng Y, Xiao B (2008) Technical note – a risk-sensitive model for managing perishable products. *Oper Res* 56(5):1305–1311
- Geraghty MK, Johnson E (1997) Revenue management saves national car rental. *Interfaces* 27(1):107–127
- Gönsch J, Klein R, Neugebauer M, Steinhardt C (2013) Dynamic pricing with strategic customers. *J of Bus Econ (JBE|ZfB)* 83(5):505–549
- Hintsches A, Spengler TS, Volling T, Wittek K, Priegnitz G (2010) Revenue management in make-to-order manufacturing: Case study of capacity control at ThyssenKrupp VDM. *BuR Bus Res* 3(2):173–190

- Hood ISA (2000) MERLIN: Model to evaluate revenue and loadings for intercity. In: Ingold A, McMahon-Beattie U, Yeoman I (eds) Yield management: Strategies for the service industries, 2nd edn. Thomson, London, pp 98–107
- Hoseason J, Johns N (1998) The numbers game: The role of yield management in the tour operations industry. *Prog in Tour and Hosp Res* 4(3):197–206
- Huang K, Chang KC (2011) A model for airline seat control considering revenue uncertainty and risk. *J of Revenue and Pricing Manag* 10(2):161–171
- 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
- Klein, JK (2000) Yield Management als Methode zur ertragsorientierten Kapazitätsnutzung bei Reiseveranstaltern – eine empirische Untersuchung. *Tour J* 4(3):283–307
- König M, Meissner J (2009a) Risk management policies for dynamic capacity control. Working paper, University of Lancaster
- König M, Meissner J (2009b) Risk minimizing strategies for revenue management with target values. Working paper, University of Lancaster
- König M, Meissner J (2010) Value-at-risk optimal policies for revenue management problems. Working paper, University of Lancaster
- Lancaster J (2003) The financial risk of airline revenue management. *J of Revenue and Pricing Manag* 2(2):158–165
- Lautenbacher CJ, Stidham SJ (1999) The underlying markov decision process in the single-leg airline yield-management problem. *Transp Sci* 33(2):136–146
- Lee T, Hersh M (1993) A model for dynamic airline seat inventory control with multiple seat bookings. *Transp Sci* 27(3):252–265
- Levin Y, McGill JI, Nediak M (2008) Risk in revenue management and dynamic pricing. *Oper Res* 56(2):326–343
- McGill JI, van Ryzin GJ (1999) Revenue management: Research overview and prospects. *Transp Sci* 33(2):233–256
- Pflug GC (2001) Some remarks on the value-at-risk and the conditional value-at-risk. In: Uryasev SP (ed) *Nonconvex optimization and its applications (vol 49), Probabilistic constrained optimization: Methodology and applications*. Kluwer, Dordrecht, pp 272–281
- Pflug GC, Pichler A (2012) On dynamic decomposition of multistage stochastic programs. Working paper, University of Vienna
- Phillips RL (2005) *Pricing and revenue optimization*. Stanford University Press, Stanford
- Rehkopf S (2006) *Revenue Management-Konzepte zur Auftragsannahme bei kundenindividueller Produktion: Am Beispiel der Eisen und Stahl erzeugenden Industrie*. Deutscher Universitäts-Verlag, Wiesbaden

- Singh P (2011) Risk in revenue management. *Inform Anal Mag*, May/June 2011:16–19
- Spengler TS, Rehkopf S, Volling T (2007) Revenue management in make-to-order manufacturing – an application to the iron and steel industry. *OR Spectr* 29(1):157–171
- Subramanian J, Stidham S, Lautenbacher CJ (1999) Airline yield management with overbooking, cancellations, and no-shows. *Transp Sci* 33(2):147–167
- Talluri KT, van Ryzin GJ (2004) *The theory and practice of revenue management*. Springer, New-York
- Tscheulin DK, Lindenmeier J (2003) Yield-Management – Ein State-of-the-Art. *Z für Betr Wirtsch* 73(6):629–662
- Volling T, Akyol DE, Wittek K, Spengler TS (2012) A two-stage bid-price control for make-to-order revenue management. *Comp & Oper Res* 39(5):1021–1032
- Weatherford LR (2004) EMSR versus EMSU: Revenue or utility? *J of Revenue and Pricing Manag* 3(3):277–284
- Weatherford LR, Bodily SE (1992) A taxonomy and research overview of perishable-asset revenue management: Yield management, overbooking, and pricing. *Oper Res* 40(5):831–844
- Wiggershaus J (2008) *Ein wertorientiertes Auftragsannahmeverfahren für das Revenue Management in der Sachgüterindustrie*. Verlag Dr. Kovač, Hamburg
- Wozabal D (2012) Value-at-risk optimization using the difference of convex algorithm. *OR Spectr* 34(4):861–883
- Xylander JK (2003) *Kapazitätsmanagement bei Reiseveranstaltern: Entscheidungsmodelle zur Kontingentierung im Yield Management*. Deutscher Universitäts-Verlag, Wiesbaden
- You PS (2008) An efficient computational approach for railway booking problems. *Eur J of Oper Res* 185(2):811–824
- Zhuang W, Li MZF (2011) Revenue optimization of risk-averse managers with atemporal utility. *J of Revenue and Pricing Manag* 10(5):424–437