



Lancaster University
MANAGEMENT SCHOOL

Lancaster University Management School
Working Paper
2009/038

Risk Management Policies for Dynamic Capacity Control

M Koenig and Joern Meissner

The Department of Management Science
Lancaster University Management School
Lancaster LA1 4YX
UK

© M Koenig and Joern Meissner
All rights reserved. Short sections of text, not to exceed
two paragraphs, may be quoted without explicit permission,
provided that full acknowledgement is given.

The LUMS Working Papers series can be accessed at <http://www.lums.lancs.ac.uk/publications/>
LUMS home page: <http://www.lums.lancs.ac.uk/>

Risk Management Policies for Dynamic Capacity Control

Matthias Koenig
Lancaster University
Management School, UK
m.koenig@lancaster.ac.uk

Joern Meissner
Lancaster University
Management School, UK
joe@meiss.com

December 6, 2009

Abstract

Consider a dynamic decision making model under risk with a fixed planning horizon, namely the dynamic capacity control model. The model describes a firm, operating in a monopolistic setting and selling a range of products consuming a single resource. Demand for each product is time-dependent and modeled by a random variable. The firm controls the revenue stream by allowing or denying customer requests for product classes. We investigate risk-sensitive policies in this setting, for which risk concerns are important for many non-repetitive events and short-time considerations.

Analyzing several numerically risk-averse capacity control policies in terms of standard deviation and conditional-value-at-risk, our results show that only a slight modification of the risk-neutral solution is needed to apply a risk-averse policy. In particular, risk-averse policies which decision rules are functions depending only on the marginal values of the risk-neutral policy perform well. The risk sensitivity of a policy only depends on the current state but it does not matter whether risk-neutral or risk-averse decisions led to the state. From a practical perspective, the advantage is that a decision maker does not need to compute any risk-averse dynamic program. Risk sensitivity can be easily achieved by implementing risk-averse functional decision rules based on a risk-neutral solution.

Keywords: dynamic decisions, capacity control, revenue management, risk

Please reference this paper as:

Matthias Koenig, Joern Meissner. Risk Management Policies for Dynamic Capacity Control. Working Paper (available at <http://www.meiss.com>), Lancaster University Management School, 2009.

BIB_T_EX and plain text references are available for download here:

<http://www.meiss.com/en/publications/risk-management-capacity-control.html>

1 Introduction

Consider a dynamic decision model under risk for capacity control with a given planning time horizon. The decision maker acts on previous gained information up to a distinct time period and estimations for future time periods. This kind of dynamic decision making under risk is often modeled by dynamic programming formulations. Despite some known limitations of expected utility theory, as discussed by Schoemaker (1982), the expected utility approach is often used with dynamic programming for risk considerations. Then, dynamic programming uses an utility function as an objective function, and time preferences can be included by a discount factor. The books of Chavas (2004) and Bertsekas (2005) include a description of this approach from a general perspective.

The considered capacity control model is typical, for example, in the area of revenue management, which use is common in industries such as airlines, hotels or rental cars: a firm operates in a monopolistic setting offering multiple products consuming a single resource. The firm owns a fixed capacity of the resource which has to be sold over a finite horizon. The objective of the firm is to find a policy in order to optimize total revenue by allocating capacity to different classes of demand. Usually, a risk-neutral optimization objective is sufficient for revenue management problems due to the long-term average effect in situations with repeating problem processes.

There are, however, many situations when the number of reiterations is too small (e.g. Levin et al. (2008) mention an event promoter) or when constraints on working capital or revenue streams force use of a dynamic decision model with consideration of risk.

The dynamic capacity control model introduced by Lee and Hersh (1993) is a standard revenue management model. The dynamic capacity control model is usually stated as dynamic programming formulation. The original approach presents a risk-neutral policy, whereas recent approaches by Barz and Waldmann (2007) and Huang and Chang (2009) propose risk-averse policies for this model.

Barz and Waldmann (2007) analyze the dynamic capacity control model under constant absolute risk-aversion using an exponential utility as objective function in the dynamic programming recursion. Huang and Chang (2009) present a policy which includes a discount factor not in the objective but in the decision function. This discount factor actually de-

termines a risk premium for certainty of earning revenue now, instead under uncertainty later. This kind of risk premium is more easily communicated to practitioners than in the case of the exponential utility function, where the computation of the risk premium requires certain knowledge about the distribution of the demand function. Huang and Chang (2009) also propose a policy considering the selling history and conduct an extensive analysis for risk-aversion and compares standard deviations and Sharpe ratios of risk-neutral and risk-averse policies.

Our objective is the evaluation of a set of control policies under risk considerations. To this end, we conduct an analysis of the policies by numerical experiments and look at risk measures in terms of volatility by the standard deviation and in terms of downside risk by the conditional-value-of-risk (CVaR). We extend the analysis of Huang and Chang (2009) and propose improved policies which are also easily implemented in practice. Furthermore, we introduce a new straight-forward policy which provides solid results for moderate levels of risk-aversion.

In particular, we demonstrate that no extra dynamic programming recursions are required for implementing decision rules for risk-sensitive policies. The risk-averse decision can be applied directly using the results of the risk-neutral case. In revenue management terms, it is sufficient to use decision rules directly with the marginal seat values of the dynamic programming solution of the risk-neutral case. Advantages include less requirements on computing a risk-averse solution and an easy and understandable way of implementing such a solution for practitioners.

The remainder of this paper is as follows. Section 2 gives a summary of related work about risk considerations in a revenue management context. We describe the dynamic capacity control model and risk-neutral as several risk-averse policies in Section 3. In Section 4, the settings of the numerical experiments and the obtained risk measures evaluating the policies are presented. Finally, we summarize and conclude this paper in Section 5.

2 Related Work

A general but comprehensive coverage of revenue management is provided by Talluri and van Ryzin (2005) for risk-neutral decision makers. Chiang et al. (2007) give an extensive literature overview about the field.

Describing the first revenue management model incorporating risk, the model of Feng and Xiao (1999) considers a single-resource problem with two given prices and allows only one price change. They define risk by sales variance as a result of price changes. Their objective function combines expected revenue and a weighted penalty function for sales variance. The weight determines the level of risk-aversion. Although their model is limited, the derived result is quite intuitive: risk-averse firms switch to a lower price sooner than risk-neutral ones. This coincides with the risk-averse policies described in Section 3, where firms prefer to accept revenue sooner rather than later.

Lancaster (2003) looks at the risk issues in airline revenue management from a practical perspective. He illustrates the vulnerability of revenue management systems by analyzing the volatility of historical data of revenue per available seat mile. He runs several simulations which highlight the importance of risk considerations under thin profit margins and high uncertainty. Therefore, he recommends a relative revenue per available seat mile at risk metric which integrates risk measurement with the value-at-risk metric. This metric is the expected maximum of underperformance over a time horizon at a choice confidence level. To compare different revenue management strategies, he proposes the use of the Sharpe ratio instead of direct dual objective optimization. This is computationally impractical as revenue distributions are acquired by history or simulations. The arguments of Lancaster (2003) also hold for our approach of comparing risk measures of different policies for dynamic capacity control.

Risk sensitivity is incorporated by Levin et al. (2008) into a dynamic pricing model of perishable products. Their objective function consists of maximum expected revenue constrained by a desired minimum level of revenue with minimum acceptable probability. This constraint is a value-at-risk formulation, and their approach corresponds with maximizing expected return subject to small disaster probability. Risk-aversion is introduced in the objective function as a penalty term reflecting the probability that total revenues fall below a

certain level. Thus, the underlying utility function at every point in time is piece-wise linear and discontinuous at the point of the desired revenue level.

Discussing risk modeling for traffic and revenue management in networks, Mitra and Wang (2005) analyze mean-variance, mean-standard-deviation and mean-conditional-value-at-risk for formulation of the objective function, finally selecting standard deviation as the risk index. The impact of several levels of risk-aversion is demonstrated by an efficient frontier for a truncated Gaussian demand distribution.

Koenig and Meissner (2010) compare expected revenue and risk in terms of standard deviation and conditional-value-at-risk of pricing policies. A list pricing policy, following capacity control and a dynamic pricing policy adjusting steadily prices, is analyzed under consideration of the cost of price changes. They show under which circumstances a policy might be in advantage of the other by numerical experiments.

Robust optimization as means for maximizing over a set of worst case outcomes under guaranteed feasibility has been covered by various authors in a revenue management context. The worst outcomes are all smallest revenue under feasible worst-case demand realizations. In this way, the works of Thiele (2006), Perakis and Roels (2007), Lai and Ng (2005), Lim and Shanthikumar (2007) and Lim et al. (2008) address the problem of uncertainty about the demand function by robust optimization. Lai and Ng (2005) set up a robust optimization solution for hotel revenue as a expected revenue versus mean absolute deviation trade-off formulation, while Lim and Shanthikumar (2007) show that the robust pricing problem is equivalent with a single-product revenue management problem with an exponential utility function without model uncertainty.

Using expected utility theory in the revenue management context is endorsed by Weatherford (2004). He discusses the assumption of risk-neutrality for a standard revenue management algorithm and concludes that optimizing expected utility instead expected revenue is a suitable risk-averse strategy. In particular, he proposes the expected marginal seat utility (EMSU) heuristic for accounting for risk-aversion instead of the expected marginal seat revenue model (EMSR), the standard algorithms introduced by Beloba (1989).

Thus, the work of Barz and Waldmann (2007) and Feng and Xiao (2008) can be considered as following the same path, employing expected utility theory for revenue management pointed out by Weatherford (2004). In order to reflect a decision maker's risk sensitivity,

both papers propose the use of an exponential utility function. Feng and Xiao (2008) show the closed form solution in this case. Barz and Waldmann (2007) considers static and dynamic capacity control models separately. In such a setting, a risk-averse policy will accept lower prices earlier in time and remaining capacity. Barz (2007) also analyze capacity control models by additive time-separable and atemporal utility function.

Huang and Chang (2009) modify the decision function for the dynamic capacity control model in order to become more risk-sensitive in terms of mean versus standard deviation. The modified decision function relaxes the optimal condition using a discounted marginal seat value. Moreover, they propose a time- and seat-dependent compromise policy which uses a hyperbolic tangent function in order to control the discount factor regarding the number of remaining seats. As proposed by Lancaster (2003), the Sharpe ratio is here applied to rank policies regarding revenue per unit of risk defined by standard deviation.

We examine the dynamic capacity policies of Barz and Waldmann (2007) and Huang and Chang (2009) in Section 3. Compared with the risk neutral policy of Lee and Hersh (1993), the risk-averse policies have in common that the acceptance of earlier certain revenue is preferred to later, possibly higher, revenue. The risk-averse policies, however, differ in when exactly to accept, depending on remaining capacity and time.

3 Description of Model and Policies

The capacity control model stated by Lee and Hersh (1993) is originally stated in the context of airline revenue management. In the risk-neutral case, the aim of the airline is to derive an optimal policy for maximization of expected revenue over a booking period under assumed demand probability for fare classes.

The booking period for a single-leg flight is divided into N decision periods such that, at worst, one request arrives per period. The number of booking classes is k and each accepted booking request results in revenue F_i and $F_1 > F_2 > \dots > F_k$. The seat request probability is based on a Poisson arrival process and the probability of a request of fare class i in decision period n is p_{in} . The probability for no booking request at all is $p_{0n} = 1 - \sum_{i=1}^k p_{in}$. The initial capacity of available seats is given by C and remaining seats are denoted by $c \leq C$.

3.1 Risk-Neutral Policy

Similar to Barz and Waldmann (2007), we describe the model by a finite-state Markov decision problem. A Markov decision problem is described by state space, action set, decision epochs, rewards and transition probabilities. In our case, these are:

- State space $S = \{0, 1, \dots, C\} \times \{0, 1, \dots, k\}$ where the first element stands for the remaining seat capacity and the second element for the fare class, with artificial fare class 0 with fare $F_0 = 0$. A state $(c, i) \in S$ says that as c seats are remaining, we have a request for fare class i .
- Action set $A(c, i) = \{0, 1\}, \forall (c, i) \in S | i > 0$ and $A(c, 0) = \{0\}$ represents the 'reject' and 'accept' decision for a given state.
- Decision epochs correspond to the time periods: $T = \{0, 1, \dots, N\}$ with time $n \in T$ is the remaining time until end of the period, the departure of the flight.
- Rewards $r_n(s, a)$ are defined for $s \in S$ and $a \in A$ by $r_n((c, i), a) = aF_i$ for $n, c > 0$.
- Transition probabilities are defined for $(c, i), (c - a, j) \in S$ and $a \in A$ by $q_n((c - a, j) | (c, i), a) = p_{jn}$ for $n = N, N - 1, \dots, 0$ and otherwise $q_n = 0$.

A sequence of decision rules $a_n = d_n(c_n, i_n)$ determines a policy $\pi = \{d_n, d_{n-1}, \dots, d_1\}$ of actions if a booking request is accepted or rejected in state (c_n, i_n) . With this setting, the expected revenue for a particular policy π starting with capacity c and request i is

$$V_N^\pi(c, i) = \mathbb{E}_{c,i}^\pi \left[\sum_{n=1}^N r_n((c_n, i_n), d_n(c_n, i_n)) + r_0(c_0, i_0) \right].$$

The maximal expected revenue may be computed by the Bellman equation for this problem:

$$V_n^*(c, i) = \max_{a \in A(c,i)} \left\{ aF_i + \sum_{j=0}^k p_{jn} V_{n-1}^*(c - a, j) \right\}, \quad (1)$$

with $V_0 = r_0 = 0$. An associated policy π^* is optimal and can be described by a decision rule $d_n^*(c, i)$ as follows (see Lee and Hersh, 1993)

$$d_n^* = \begin{cases} 1, & F_i > \sum_{i=0}^k p_{in}(V_{n-1}^*(c, i) - V_{n-1}^*(c-1, i)) \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

This policy can be explained as accepting a request only, if the fare F_i of the request is greater than the expected marginal seat value $\delta_{V_{n-1}}(c, i) = \sum_{i=0}^k p_{in}(V_{n-1}(c, i) - V_{n-1}(c-1, i))$.

3.2 Risk-Sensitive Policies

In this section, we present several risk-sensitive policies. The first policy was published by Barz and Waldmann (2007) and uses an exponential utility function. The two further policies were presented by Huang and Chang (2009). One discounts the marginal seat value in the decision rule, and another uses a time- and seat-dependent policy. These policies recursively compute dynamic programming solutions, as the proposed decision rules depend on the marginal seat values computed within.

We propose that it is not required to compute extra dynamic programming solutions for every risk-averse decision. We further demonstrate that the accumulated risk-averse development of the marginal seat values is not necessary in order to apply risk-aversion. Instead, we introduce risk-averse policies which decision rules depend only on the marginal seat values of the risk-neutral solution.

Finally, we show a policy which takes into consideration the state of the ratio between remaining capacity and time periods. This policy can be considered as a simpler but similar approach than the time- and seat-dependent policy of Huang and Chang (2009). We compare all policies in the next section numerically.

3.2.1 Exponential Utility Function

Barz and Waldmann (2007) introduce an exponential utility function and employ results of Howard and Matheson (1972) in order to derive an optimal policy for this approach. An

exponential utility function has the form $u_\gamma(x) = -\exp(-\gamma x)$ with positive parameter γ determining the level of risk-aversion.

Thus, the Markov decision process is changed and the expected value of a policy $\pi^\gamma = \{d_n^\gamma, d_{n-1}^\gamma, \dots, d_1^\gamma\}$ is now

$$V_N^{\pi^\gamma} = \mathbb{E}_{(c,i)}^{\pi^\gamma} \left[-\exp \left(-\gamma \sum_{n=1}^N r_n((c_n, i_n), d_n(c_n, i_n)) + r_0(c_0, i_0) \right) \right].$$

The computation of the maximal expected exponential utility leads to

$$V_n^{*\gamma}(c, i) = \max_{a \in A(c, i)} \left\{ \exp(-\gamma a F_i) \cdot \sum_{j=0}^k p_{jn} V_{n-1}^{*\gamma}(c - a, j), \right\} \quad (3)$$

and $V_0^{*\gamma} = -\exp(-\gamma V_0^*(c, i)) \forall (c, i) \in S$. Corresponding optimal policies are $\pi^{*\gamma} = \{d_N^{*\gamma}, d_{N-1}^{*\gamma}, \dots, d_1^{*\gamma}\}$. The γ -optimal policy can be derived and results in

$$d_n^{*\gamma} = \begin{cases} 1, & \exp(-\gamma F_i) < \frac{\sum_{i=0}^k p_{in} V_{n-1}^{*\gamma}(c, i)}{\sum_{i=0}^k p_{in} V_{n-1}^{*\gamma}(c-1, i)} \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

A request is accepted if its utility is lower than the expected utility gain of an additional seat.

3.2.2 Discounted Marginal Seat Value in Dynamic Programming Recursion

One of the policies which Huang and Chang (2009) propose is the relaxation of optimality for a more risk-sensitive policy. They show in a numerical experiment the behavior of this policy in terms of average and standard deviation. This policy discounts the marginal seat value. Its value function is

$$V_n^\beta = \begin{cases} F_i + \sum_{i=0}^k p_{in} V_{n-1}^\beta(c-1, i) & F_i \geq \beta \cdot \delta_{V_{n-1}^\beta}(c, i) \\ \sum_{i=0}^k p_{in} V_{n-1}^\beta(c, i) & \text{otherwise,} \end{cases} \quad (5)$$

and the according policy π^β has decision rules

$$d_n^\beta(c_n, i_n) = \begin{cases} 1, & F_i \geq \beta \cdot \delta_{V_{n-1}^\beta}(c, i) \\ 0, & \text{otherwise,} \end{cases} \quad (6)$$

where β is a discount factor which should be in the range $[0, \dots, 1]$.

3.2.3 Discounted Marginal Seat Value within Risk-Neutral Solution

Based on the previous idea, we propose using its decision rules but directly on the marginal seat values of the risk-neutral solution. Instead of using the marginal seat value δ_{V^β} determined by the recursive value function of Equation 5, we work directly with the marginal seat value δ_{V^*} with a prior computed risk-neutral value function (Equation 1). Thus, the decision rules do not depend on previous 'discounted' decisions. We require only the risk-neutral solution and can transform this into policy $\pi^{\beta, \nu}$ which uses

$$d_n^{\beta, \nu}(c_n, i_n) = \begin{cases} 1, & F_i \geq \beta \cdot \delta_{V_{n-1}^*}(c, i) \\ 0, & \text{otherwise,} \end{cases} \quad (7)$$

where β is again a discount factor. The value function V_n^β would be transformed similarly to obtain $V_n^{\beta, \nu}$. Note that the decisions do not require pre-computing any risk-averse solution but are based on the risk-neutral one. However, the (c, i) path may be affected by previous decisions; whether they were risk-neutral or risk-averse based need not be known at the current state.

3.2.4 Selling-Rate Dependent Decisions

We consider two selling-rate dependent decisions. Basically, this kind of policy increases the level of risk-aversion if less than the expected number seats have been sold up the current time period.

We start with the time- and seat-dependent compromise policy of Huang and Chang (2009). It uses a hyperbolic tangent function and two variable parameter κ_1 and κ_2 which determine the level of risk-sensitive behavior in dependence on the number of remaining

seats before departure. The discount factor $\beta_n(c)$ is computed as

$$\beta_n^{k_1, k_2}(c) = \frac{1}{2} \left[\tanh \left(\kappa_1 \left(C \frac{\sum_{m=1}^n \sum_{i=1}^k F_i p_{in}}{\sum_{m=1}^N \sum_{i=1}^k F_i p_{in}} + \kappa_2 - c \right) \right) + 1 \right].$$

This time- and seat-dependent factor can be used in either of the policies π^β and $\pi^{\beta, \nu}$ which we will denote with π^{β_κ} and $\pi^{\beta_\kappa, \nu}$.

Based on this policy, we introduce a further policy which takes also into account the selling of seats per time period. We define a function which serves as reference if sales of seats develop as expected 'on-track'. This function is integrated in an indicator function which helps us to find out if the sales rate diverges from the track:

$$\mathbb{1}_n(c) = \begin{cases} 1, & c > C \frac{\sum_{m=1}^n \sum_{i=1}^k p_{in}}{\sum_{m=1}^N \sum_{i=1}^k p_{in}} \\ 0, & \text{otherwise.} \end{cases} \quad (8)$$

The indicator function is embedded into the decision rules in order to switch on and off a discount factor according to the current state of the selling-rate. We define policy $\pi^{\beta, \mathbb{1}}$ by the decision rules

$$d_n^{\beta, \mathbb{1}}(c_n, i_n) = \begin{cases} 1, & F_i \geq \beta^{\mathbb{1}_n(c)} \cdot \delta_{V_{n-1}^*}(c, i) \\ 0, & \text{otherwise.} \end{cases} \quad (9)$$

This policy $\pi^{\beta, \mathbb{1}}$ is distinguished from π^{β_κ} policies by a hard on and off switch for comparing with or without discount in the decision rules. Again, we can apply this policy using the marginal seat values inside a risk-sensitive recursive formulation or using the marginal seat values of the risk-neutral solution. We consider only the latter one already denoted $\pi^{\beta, \mathbb{1}}$.

4 Numerical Simulation and Results

The described policies are analyzed in a numerical simulation for the purpose of comparing their performance regarding revenue and risk. We follow the setup used by other authors in order to allow better comparisons. We start with a comparison of the policies using the recursive risk-sensitive marginal seat values and the risk-neutral marginal seat values. As

the results will show, the results are nearly identical. Therefore, we will continue with the $\pi^{\cdot, \nu}$ policies and compare them among themselves and with the exponential utility function approach.

4.1 Risk Measures

We evaluate the risk involved in the policies in terms of standard deviation as the measure for volatility and conditional-value-at-risk as the measure for downside risk.

4.1.1 Standard Deviation

Although standard deviation measures the dispersion of a random variable, it is often interpreted as risk measure. It gives information about the level of variation of a random variable around its expected value.

For a random variable X , the standard deviation is the square root of the moment-based measure variance σ^2 ; it is defined as $\sigma(X) = \sqrt{E(X^2) - E(X)^2}$, where E denotes the expected value of a random variable.

Standard deviation can be best considered as a volatility measure as it measures values which are better than and worse than expected value. Therefore, it does not represent only the downside (worse than expected) but also the upside (better than) return of an applied policy. For this reason, we consider downside measures, such as value-at-risk and conditional-value-at-risk, as more meaningful as risk measure.

4.1.2 Value-at-Risk and Conditional-Value-at-Risk

Value-at-risk (VaR) has become very common risk measure in the financial industry, where it originated. For a given particular confidence level and time horizon, it measures the maximum expected loss on a portfolio of assets. The investor chooses the time horizon and a confidence level. Common confidence levels are 95% or 99%, and VaR helps to estimate the maximum loss of the portfolio in 95% or 99% of cases. VaR can also be a suitable risk measure for other industrial sectors as it is actually a percentile of a random variable that presents a distribution of returns.

Let the confidence level be denoted by $1 - \alpha$ and the random variable X represent earnings. We assume here that a lower $x \in X$ means greater loss (or less return). VaR can

be defined by the α -quantile of X with distribution function P and cumulative distribution function F_X as

$$VaR_\alpha(X) = \inf\{x : P(X \leq x) \geq \alpha\} = \inf\{x : F_X(x) \geq \alpha\}.$$

A disadvantage of VaR is that it does not reveal anything about the distribution of X below the VaR of the particular confidence level $1 - \alpha$. It is not a coherent risk measure. Thus, conditional-value-at-risk (CVaR), also known as expected shortfall, is often preferred as a risk measure.

CVaR does not have the deficiencies of VaR and can be understood as the expected value given the return is less than or equal to the VaR value. Conditional-value-at-risk can be defined as:

$$CVaR_\alpha(X) = E(X|X \leq VaR_\alpha(X)).$$

We follow the arguments by Luciano et al. (2003); Lancaster (2003) and Ahmed et al. (2006) that VaR and CVaR can be useful not only in the context of investments, but also for other applications. We consider CVaR as a useful measure for revenue management policies as it provides information about the downside of achieved revenue for a given confidence and a time horizon.

4.2 Simulation Setup

The setup is taken from Lee and Hersh (1993)'s example. Four booking classes are considered with fares $F_1 = 200, F_2 = 150, F_3 = 120, F_4 = 80$. The number of time periods to departure is $N = 30$ and the initial capacity of seats is $C = 10$. The probabilities of requests regarding fare class and time period are listed in Table 1. For every experiment, we run 10,000 sample runs with average revenue, standard deviation and conditional-value-of-risk. The random arrivals are simulated in a Monte Carlo approach with the underlying probability distribution. The same random data is used for each comparison between a set of policies as shown in each figure later on.

i	F_i	$1 \leq n \leq 4$	$5 \leq n \leq 11$	$12 \leq n \leq 18$	$19 \leq n \leq 25$	$26 \leq n \leq 30$
1	200	0.15	0.14	0.10	0.06	0.08
2	150	0.15	0.14	0.10	0.06	0.08
3	120	0	0.16	0.10	0.14	0.14
4	80	0	0.16	0.10	0.14	0.14

Table 1: Fares and request probabilities for fare class i and time period n .

4.3 Results


In order to illustrate the effect of a risk-averse versus a risk-neutral policy, we present in Figure 4.3 the protection levels obtained by the different policies. The four subfigures depict the protection levels for a given time period and remaining inventory of seats. The numbers in the matrix are to be interpreted as the lowest class of which requests are to be accepted, e.g. a 'two' means that only request for the first and second class are accepted. The ordinate shows the remaining seats or inventory. The abscissae displays the remaining time periods before departure. Time period zero of departure is on the left hand side.

The visualization of the protection levels gives an impression how the risk-aversion influences the optimal risk-neutral policy. The acceptance of booking requests of lower classes shifts to earlier time periods. This is observable as all risk-averse policies open earlier all four classes according remaining inventory as well as remaining time periods. Figure 4.3(a), (b) and (c) also illustrate differences between policies and level of risk-aversion. As the protection levels seem to shift more to the right hand side and to the bottom, the three figures show increasing levels of risk-aversion.


Figure 4.3(d) displays a seat and inventory dependent policy. This becomes good visible as the protection levels of the right hand and bottom side are similar to the risk-neutral case. However, the risk-sensitiveness is observable if inventory is high and remaining time periods are low.

4.3.1 Comparison of Decision Rules Dependent on Marginal Seat Values of Risk-Neutral and Risk-Averse Recursive Solution


The differences between applying the decision rules on basis of the marginal seat values of the risk-neutral solution and on the risk-sensitive recursive solutions are analyzed with the policies π^β and $\pi^{\beta,\nu}$ first. Both use a discount factor in order to relax the decision if




(a)



(b)



(c)



(d)

Figure 1: Protection levels generated by the different policies: π^* (a), $\pi^{*,\gamma=0.005}$ (b), $\pi^{\beta=0.8,\nu}$ (c), and $\pi^{\beta=0.8,1}$.

an arrival should be accepted or rejected. Figure 2 shows the averaged values of revenue versus standard deviation and of revenue versus CVaR. The policy $\pi^{\beta,\nu}$ slightly outperforms policy π^β in both evaluated measures, though the difference is not remarkable at all. This is recognizable as the graph of policy $\pi^{\beta,\nu}$ embeds the graph of policy π^β .


The policies applying selling-rate decision rules are compared in Figure 3. It illustrates that the differences between the policies $\pi^{\beta_{\kappa_1,\kappa_2}}$ and $\pi^{\beta_{\kappa_1,\kappa_2},\nu}$ are also negligible. There is little visible difference between the graphs.

4.3.2 Comparison of Discounted Marginal Seat Value and Selling-Rate Dependent Policies


We compare the discounted marginal seat value using policy $\pi^{\beta,\nu}$ with the two selling-rate dependent policies $\pi^{\beta_{\kappa_1,\kappa_2},\nu}$ and $\pi^{\beta,\mathbb{1},\nu}$. Here, we use the $\pi^{\beta,\nu}$ as representative of both policies using the discounted marginal seat value because it performed at least as the other. We use $\pi^{\beta_{\kappa_1,\kappa_2},\nu}$ instead of $\pi^{\beta_{\kappa_1,\kappa_2}}$, but as both showed the same performance, they are exchangeable.

Figure 4(a) shows that the policies $\pi^{\beta,\nu}$ and $\pi^{\beta_{\kappa_1,\kappa_2},\nu}$ have similar volatility. There is no visible difference between $\pi^{\beta_{\kappa_1,\kappa_2},\nu}$ compared with $\pi^{\beta,\nu}$ in terms of a lower standard deviation. The second selling-rate dependent policy $\pi^{\beta,\mathbb{1},\nu}$ that uses an indicator function has a higher volatility. This volatility becomes greater with increasing level of risk-aversion β .

The CVaR measure reveals different results in terms of downside risk as shown in Figure 4(b). The policy $\pi^{\beta,\nu}$ achieves less revenue in the worst 5% of cases than both selling-rate dependent policies, except for a high level for risk-aversion. This can be seen at the lower right hand corner of the figure where the graph of policy $\pi^{\beta,\mathbb{1},\nu}$ decreases and crosses policy $\pi^{\beta,\nu}$. Further, policy $\pi^{\beta_{\kappa_1,\kappa_2},\nu}$ performs well here for all level of risk sensitivity. It outperforms policy $\pi^{\beta,\nu}$. On the other hand, policy $\pi^{\beta,\mathbb{1},\nu}$ displays the same low risk as the other selling-rate dependent policy for moderate risk-sensitivity up to a risk-aversion level of $\beta = 0.6$. Then it drops down significantly. This drop coincides with its behavior for the standard deviation. However, as a policy which just switch on or off risk-aversion, its results are very good for moderate levels of risk-aversion.




(a)




(b)

Figure 2: Comparison of the polices π^β and $\pi^{\beta,v}$ using revenue vs. standard deviation (a) and revenue vs. conditional-value-of-risk (b) with $\beta \in [0.8, \dots, 1.0]$ using a step size of 0.025.




(a)




(b)

Figure 3: Comparison of the policies $\pi^{\beta_{\kappa_1, \kappa_2}}$ and $\pi^{\beta_{\kappa_1, \kappa_2, \nu}}$ using revenue vs. standard deviation (a) and revenue vs. conditional-value-of-risk (b). The value of κ_1 is fixed at 0.3 and κ_2 goes from 1.5 to 5.0 with step size of 0.5.



(a)



(b)

Figure 4: Comparison of the policies $\pi^{\beta, v}$, $\pi^{\beta_{\kappa_1, \kappa_2}, v}$ and $\pi^{\beta, \mathbb{1}, v}$ using revenue vs. standard deviation (a) and revenue vs. conditional-value-of-risk (b). Policy $\pi^{\beta, v}$ uses $\beta \in [0.80, \dots, 0.95]$ with step size 0.05. Policy $\pi^{\beta_{\kappa_1, \kappa_2}, v}$ has κ_1 fixed at 0.3 and κ_2 from 1.5 to 5.0 with step size 0.5. Policy $\pi^{\beta, \mathbb{1}, v}$ uses $\beta \in [0.5, \dots, 0.95]$ with step size 0.05.

4.3.3 Comparison of Exponential Utility and Selling-Rate Dependent Policies


Finally, we compare the selling-rate dependent policies with the policy using the exponential utility function. This is demonstrated in Figure 5. It is notable that there is a very similar behavior between policies $\pi^{\beta_{\kappa_1, \kappa_2, \nu}}$ and $\pi^{*\gamma}$. Their results differ slightly only for a high level of risk-aversion. This is observable in Figure 5(b) on the lower right hand side. Policy $\pi^{\beta, \mathbb{1}, \nu}$ can only follow the other policies up to a certain level of risk-sensitivity as already mentioned above.

5 Conclusions


Several risk-averse policies for the dynamic capacity control problem were compared regarding revenue versus standard deviation and revenue versus conditional-value-at-risk. We have shown that there are only small differences between the risk-sensitive policies employing a discounted decision rule. The results in terms of the proposed risk measures are similar and do not depend on the computation of a complete dynamic programming solution for each level of risk-aversion. We have the same situation for the two comparable policies which use a selling-rate dependent policy employing a hyperbolic tangent function. In both cases, there is no advantage of computing a complete dynamic programming solution other than the risk-neutral one.

Furthermore, we presented a new selling-rate dependent policy which only (de)activates risk-aversion in dependence of the selling-rate. This policy keeps up with a previous proposed selling-rate dependent policy and an exponential utility using policy for a wide range of moderate levels of risk-sensitivity in terms of down-side risk.

We have shown that it is adequate to apply the decisions rule on the marginal seat values of the neutral solution to achieve at least similar results than policies offer that use decision rules based on marginal seat values of distinct dynamic programming solution of risk-averse policies. From a practical point of view, this has the advantage that without requiring new computation, risk-aversion can be easily integrated in the risk-neutral solutions in a readily understandable way.



(a)



(b)

Figure 5: Comparison of the policies π^{γ} , $\pi^{\beta_{\kappa_1, \kappa_2, \nu}}$ and $\pi^{\beta, 1, \nu}$ using revenue vs. standard deviation (a) and revenue vs. conditional-value-of-risk (b). Policy π^{γ} has $\gamma \in [0.001, \dots, 0.005]$ with step size 0.0005. Policy $\pi^{\beta_{\kappa_1, \kappa_2, \nu}}$ has κ_1 fixed at 0.3 and κ_2 from 1.5 to 5.0 with step size 0.5. Policy $\pi^{\beta, 1, \nu}$ uses $\beta \in [0.5, \dots, 0.95]$ with step size 0.05.

References

- Ahmed, S., Cakmak, U., and Shapiro, A. (2006). Coherent risk measures in inventory problems. *European Journal of Operational Research*, 182(1):226–238.
- Barz, C. (2007). *Risk-Averse Capacity Control in Revenue Management*. PhD thesis, University of Karlsruhe.
- Barz, C. and Waldmann, K.-H. (2007). Risk-sensitive capacity control in revenue management. *Mathematical Methods of Operations Research*, 65(3):565–579.
- Beloba, P. P. (1989). Application of an probabilistic decision model to airline seat inventory control. *Operations Research*, 37(2):183–197.
- Bertsekas, D. P. (2005). *Dynamic Programming and Optimal Control, Vol. I*. Athena Scientific, 3rd edition.
- Chavas, J.-P. (2004). *Risk Analysis in Theory and Practice*. Elsevier Academic Press.
- Chiang, W.-C., Chen, J., and Xu, X. (2007). An overview of research on revenue management: Current issues and future research. *International Journal of Revenue Management*, 1(2):97–128.
- 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.
- Howard, R. A. and Matheson, J. E. (1972). Risk-sensitive markov decision processes. *Management Science*, 18(7):356–369.
- Huang, K. and Chang, K. (2009). A simulation analysis for the airline seat control considering expected revenue and risk. *Journal of Revenue and Pricing Management*.
- Koenig, M. and Meissner, J. (2010). List pricing versus dynamic pricing: Impact on the revenue risk. Forthcoming in *European Journal of Operational Research*.

- Lai, K.-K. and Ng, W.-L. (2005). A stochastic approach to hotel revenue optimization. *Computers and Operations Research*, 32(5):1059-1072.
- Lancaster, J. (2003). The financial risk of airline revenue management. *Journal of Revenue and Pricing Management*, 2(2):158-165.
- Lee, T. and Hersh, M. (1993). A model for dynamic airline seat inventory controls 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.
- Lim, A. and Shanthikumar, J. (2007). Relative entropy, exponential utility, and robust dynamic pricing. *Operations Research*, 55(2):198-214.
- Lim, A. E. G., Shanthikumar, J. G., and Watewal, T. (2008). Robust multi-product pricing. Technical report, University of Berkeley.
- Luciano, E., Peccati, L., and Cifarelli, D. M. (2003). VaR as a risk measure for multiperiod static inventory models. *International Journal of Production Economics*, 81-82:375-284.
- Mitra, D. and Wang, Q. (2005). Stochastic traffic engineering for demand uncertainty and risk-aware network revenue management. *IEEE/ACM Transactions on Networking*, 13(2):221-233.
- Perakis, G. and Roels, G. (2007). Robust controls for network revenue management. Technical report, Decisions, Operations, and Technology Management, University of California.
- Schoemaker, P. J. H. (1982). The expected utility model: Its variants, purposes, evidence and limitations. *Journal of Economic Literature*, 20:529-563.
- Talluri, K. and van Ryzin, G. (2005). *The Theory and Practice of Revenue Management*. Springer.
- Thiele, A. (2006). Single-product pricing via robust optimization. Technical report, Lehigh University.
- Weatherford, L. (2004). EMSR versus EMSU: Revenue or utility? *Journal of Revenue and Pricing Management*, 3(3):277-284.