Risk Management Policies for Dynamic Capacity Control

Matthias Koenig[a], Joern Meissner[b]

[a]*Department of Management Science, Lancaster University Management School, UK*
[b]*Kuehne Logistics University, Hamburg, Germany*

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.

Numerically analyzing several 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. 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

---

Email addresses: matthias.koenig@fh-bielefeld.de (Matthias Koenig), joe@meiss.com (Joern Meissner)

[a]now at: University of Applied Sciences Bielefeld, Germany

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 modelled by dynamic programming formulations. Despite some known limitations of expected utility theory, as discussed by Schoemaker [1], the expected utility approach is often used with dynamic programming for risk considerations. To this end, dynamic programming uses a utility function as an objective function, and time preferences can be included by a discount factor. The books by Chavas [2] and Bertsekas [3] 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, whose use is common in industries such as airlines, hotels or rental cars, in which 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 optimise total revenue by allocating capacity to different classes of demand. Usually, a risk-neutral optimisation objective is sufficient for revenue management problems due to the long-term average effect in situations with repeating decision-making processes.

There are, however, many situations when the number of reiterations is too small (e.g., Levin et al. [4] mention an event promoter) or when constraints on working capital or revenue streams force the use of a dynamic decision model with consideration of risk. Weatherford [20] observes that analysts were uncomfortable with risk-neutral objectives and changed waiting lists recommended by their revenue management systems. This means, they applied manually their own risk-averse policy. In practice, short-time objectives of management are a motivation for risk aversion as pointed out by Feng and Xiao [22]. The authors emphasise the obvious effect that uncertainty in demand, forecast and capacity may lead to a significant difference between the realised revenue and expected revenue. The practitioners' demand for risk aversion has motivated the research of risk-averse policies and, thus, this paper, too.

Furthermore, recent approaches by Barz and Waldmann [6] and Huang and Chang [7] propose risk-averse policies for the dynamic capacity control model. This model is introduced a standard revenue management model by Lee and Hersh [5] and is originally stated as a dynamic programming formulation of a risk-neutral policy. Barz and Waldmann [6] analyse the dynamic capacity control model under constant absolute risk aversion using an exponential utility as the objective function in the dynamic programming recursion. Huang and Chang [7] present a policy which includes a discount factor not in the objective but in the decision function. This discount factor actually determines a risk premium for certainty of earning revenue now, instead of under uncertainty later. This kind of risk premium is more easily communicated to practitioners than the exponential utility function, where the computation of the risk premium requires certain knowledge about the distribution of the demand function. Huang and Chang [7] also propose a policy considering the selling history and conduct an extensive analysis for risk aversion, comparing 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 perform 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 [7] and propose improved policies which are also easily implemented in practice. Furthermore, we introduce a new straightforward policy which provides acceptable results for moderate levels of risk aversion.

The contribution of this paper is an improvement for applying risk-averse policies. Our presented approach offers advantages for a decision maker in terms of computational and memory requirements. The advantages include less requirements on computing a risk-averse solution and an easy and understandable way of implementing such a solution for practitioners. This is achieved as only one dynamic programming solution is needed for the application of policies of various levels of risk sensitivity. As decision makers often have to determine their level of risk aversion by trying out different levels in simulation, where each level requires a dynamic programming solution, they benefit from our method which requires only the risk-neutral solution for each level. Also, the risk-neutral solution can serve as a basis for applying different levels of risk sensitivity to certain instances in the same setting, i.e., when the risk level is changed from instance to instance. Additionally, we propose a policy which allows to switch risk aversion on or off depending on the current state of the selling rate. This proposed approach could be used with a de(activation) of risk aversion dependent on other possible variables, too.

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 capacity values of the dynamic programming solution of the risk-neutral case.

The remainder of this paper is as follows. Section 2 gives a summary of related work about risk considerations in revenue management context. We describe the dynamic capacity control model, risk-neutral and risk-averse policies associated with the model in Section 3. In Section 4, we present the settings of the numerical experiments and the obtained risk measures evaluating the policies. Finally, we summarise 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 [8] for risk-neutral decision makers. Chiang et al. [9] give an extensive literature overview of the field.

The first revenue management model incorporating risk, the model of Feng and Xiao [10], 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 [11] looks at the risk issues in airline revenue management from a practical perspective. He illustrates the vulnerability of revenue management systems by analysing 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 chosen confidence level. To compare different revenue management strategies, he proposes the use of the Sharpe ratio instead of direct dual objective optimisation. This is computationally impractical as revenue distributions are acquired by history or simulations. The arguments [11] for using simulations also hold for our approach of comparing risk measures of different policies for dynamic capacity control.

Risk sensitivity is incorporated by Levin et al. [4] 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 maximising expected return subject to a 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 modelling for traffic and revenue management in networks, Mitra and Wang [12] analyse 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 [13] 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, steadily adjusting prices, are analysed under consideration of the cost of price changes. They show by numerical experiments under which circumstances a policy might be more advantageous over the other.

Robust optimisation [cf. 14] as a means for maximising over a set of worst case outcomes under guaranteed feasibility has been used by various authors in a revenue management context. The worst outcomes are all the smallest revenues under feasible worst-case demand realisations. The works of Perakis and Roels [15], Thiele [16] and Lim and Shanthikumar [17] are exemplary for addressing the problem of uncertainty in the demand function by robust optimisation. Lim and Shanthikumar [17] show that the robust pricing problem is equivalent to a single-product revenue management problem with an exponential utility function without model uncertainty.

Mulvey et al. [18] propose a different approach and consider robustness of solutions in a set of scenarios. They introduce a penalty function to the objective function to achieve a tradeoff between optimality and feasibility. Following this approach, Lai and Ng [19] set up a model for hotel revenue and formulate a tradeoff between expected revenue and mean absolute deviation.

Using expected utility theory in the revenue management context is endorsed by Weatherford [20]. He discusses the assumption of risk-neutrality for a standard revenue management algorithm and concludes that optimising expected utility instead of 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 [21].

Thus, the work of Barz and Waldmann [6] and Feng and Xiao [22] can be considered as following the same path, employing expected utility theory for revenue management as pointed out by Weatherford [20]. In order to reflect a decision maker's risk sensitivity, both papers propose the use of an exponential utility function. Feng and Xiao [22] show the closed form solution in this case. Barz and Waldmann [6] consider static and dynamic capacity control models separately. In such a setting, a risk-averse policy will accept lower prices earlier in time and in higher remaining capacity. Barz [23] also analyses capacity control models by additive time-separable and atemporal utility function.

Capacity control policies for make-to-order revenue management problems are addressed by Volling et al. [24]. The authors propose a two-stage approach for capacity control in make-to-order settings. By evaluating the risk that their policy falls below the first-come-first-served policy, they present results for a risk-averse decision maker.

Huang and Chang [7] modify the decision function for the dynamic capacity control model in order to make it risk-sensitive in terms of mean versus standard deviation. The modified decision function relaxes the risk-neutral optimal decision using a discounted marginal seat value. Moreover, they propose a time- and seat-dependent policy which uses a hyperbolic tangent function in order to control the discount factor regarding the number of remaining seats. As proposed by Lancaster [11], the Sharpe ratio is here applied to rank policies regarding revenue per unit of risk, defined by the standard deviation.

We examine the dynamic capacity policies of Barz and Waldmann [6] and Huang and Chang [7] in Section 3. Compared with the risk neutral policy of Lee and Hersh [5], 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 by Lee and Hersh [5] 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 maximisation of expected revenue over a booking period under assumed demand probabilities for fare classes.

The booking period for a single-leg flight is divided into N decision periods such that, at most, one request arrives per period. The number of booking classes is k and each accepted booking request for class i 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 [6], 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 $A(c, i) = \{0, 1\}, \forall (c, i) \in S | c, i > 0$ and $A(0, i) = A(c, 0) = \{0\}$ 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 decision rule $d_n(c_n, i_n) = a_n$ determines if a booking request is accepted or rejected in the state (c_n, i_n). A policy $\pi = \{d_N, d_{N-1}, \dots d_1\}$ is built from a sequence of decision rules. 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:

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

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 5]

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

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

3.2. Risk-Sensitive Policies

In this section, we present several risk-sensitive policies. The first policy was published by Barz and Waldmann [6] and uses an exponential utility function. The two further policies were presented by Huang and Chang [7]. One discounts the marginal seat value in the decision rule, and another uses a time- and seat-dependent policy. The policies recursively compute dynamic programming solutions, as the proposed decision rules depend on the marginal seat values computed by these solutions.

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 use decision rules that 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 [7]. We compare all policies numerically in the next section.

3.2.1. Exponential Utility Function

Barz and Waldmann [6] introduce an exponential utility function and employ the results of Howard and Matheson [25] 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

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

The computation of the maximal expected exponential utility leads to

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

and $V_0^{*\gamma}(c,i) = -\exp(-\gamma V_0^*(c,i))\,\forall(c,i)\in S$. 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

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

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 [7] propose is the relaxation of optimality for a more risk-sensitive policy π^β. They show in a numerical experiment the behaviour of this policy in terms of mean and standard deviation. The policy π^β discounts the marginal seat value. Its value function is

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

and the according policy π^β has decision rules

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

where β is a discount factor which represents risk aversion and should be chosen between 0 and 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).

Equation 5 does not discount the value but relaxes the decision. Although the recursion accumulates previous decisions, the change in the value function $V_n^\beta(c,i)$ is caused by the relaxation of the current decision itself. Hence, we take a further step forward as we use only the relaxation and neglect the consequences of accumulation of previous decisions. Thus, our proposed decision rules do not depend on previous 'discounted' decisions. Using the discount factor β, the value function becomes

$$V_n^{\beta,\nu}(c,i) = \begin{cases} F_i + \sum_{j=0}^{k} p_{jn} V_{n-1}^{\beta,\nu}(c-1,j) & F_i \geq \beta \cdot \delta_{V_{n-1}^*}(c) \\ \sum_{j=0}^{k} p_{jn} V_{n-1}^{\beta,\nu}(c,j) & \text{otherwise.} \end{cases} \tag{7}$$

Though the computation of the value function is stated as a dynamic program, the application of its associated policy $\pi^{\beta,\nu}$ does not require this value function. Its decision rules depend on only the risk-neutral value function V_n^*. The decision rules of $\pi^{\beta,\nu}$ are

$$d_n^{\beta,\nu}(c,i) = \begin{cases} 1, & F_i \geq \beta \cdot \delta_{V_{n-1}^*}(c) \\ 0, & \text{otherwise.} \end{cases} \tag{8}$$

It is obvious that there is no difference between the policies π^β and $\pi^{\beta,\nu}$ for $\beta = 0$ (always accept) and $\beta = 1$ (risk-neutral). For other values of β, the decision rules of both policies differ in the cases if $\beta\delta_{V_{n-1}^*}(c) > F_i \geq \beta\delta_{V_{n-1}^\beta}(c)$. Thus, we can expect only small differences between the policies.

Again, note that the decisions $d_n^{\beta,\nu}$ 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. The risk-neutral solution is reused with different values of discount factor β in order to consider several levels of risk aversion.

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 fewer than the expected number of seats have been sold up to the current time period.

We start with the time- and seat-dependent compromise policy of Huang and Chang [7]. It uses a hyperbolic tangent function and two variable parameter κ_1 and κ_2 which determine the level of risk-sensitive behaviour that depends on the number of remaining seats before departure. The discount factor $\beta_n(c)$ is computed as

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

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 $\pi^{\beta_{\kappa_1,\kappa_2}}$ and $\pi^{\beta_{\kappa_1,\kappa_2},\nu}$.

Based on this policy, we introduce a further policy which also takes 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:

$$\mathbb{1}_n(c) = \begin{cases} 1, & c > C\frac{\sum_{m=1}^{n}\sum_{i=1}^{k} p_{im}}{\sum_{m=1}^{N}\sum_{i=1}^{k} p_{im}} \\ 0, & \text{otherwise.} \end{cases} \tag{9}$$

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

$$d_n^{\beta,\mathbb{1}}(c,i) = \begin{cases} 1, & F_i \geq \beta^{\mathbb{1}_n(c)} \cdot \delta_{V_{n-1}^*}(c) \\ 0, & \text{otherwise.} \end{cases} \tag{10}$$

The indicator function is used as exponent: $\beta^{\mathbb{1}_n(c)} = \beta$ if $\mathbb{1}_n(c) = 1$ and $\beta = 1$ otherwise. This policy $\pi^{\beta,\mathbb{1}}$ is distinguished from π^{β_κ} policies by a hard on and off switch which is used in the decision rule. The switch enables or disables the discounting of the marginal seat value when it is compared with the fare prices. 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 analysed in a numerical simulation for the purpose of comparing their performance in terms of revenue and risk. We illustrate the performance by following the setup which is used by other authors in order to allow better comparisons. Furthermore, we conducted a series of experiments with different setups which are summarised in the appendix of this paper. Too, we present the revenue distributions of a further representative experiment in the appendix. 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 when applying the policies in terms of volatility and downside risk. Standard deviation is used 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 is a measure of expected distance between values which are better or 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 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 represents 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. [26], Lancaster [11] and Ahmed et al. [27] 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

For the illustration of the result, the setup was taken from Lee and Hersh [5]'s example. Four booking classes were considered with fares $F_1 = 200, F_2 = 150, F_3 = 120, F_4 = 80$. The number of time periods to departure was $N = 30$ and the initial capacity of seats is $C = 10$. The probabilities of requests for each 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 were simulated by using a Monte Carlo approach with the probability distribution given in Table 1. The same random data was used for each policy when policies were compared. The load factor $\rho = \frac{1}{C}\sum_{n=0}^{N}\sum_{j=1}^{k} p_{jn}$ of this setup is 1.32, which means that more seat requests are expected than available seat capacity is available.

4.3. Results

In order to illustrate the effect of a risk-averse versus a risk-neutral policy, we present in Figure 1 the protection levels obtained by different

i	F_i	$1 \le n \le 4$	$5 \le n \le 11$	$12 \le n \le 18$	$19 \le n \le 25$	$26 \le n \le 30$
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.

policies. The seven 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 for which requests are to be accepted, e.g., a 'two' means that only requests for the first and second class were accepted. The ordinates show the remaining seats or inventory. The abscissae display the remaining time periods before departure. Time period zero of departure is on the left hand side.

The visualisation of the protection levels gives an impression of 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 all four classes earlier when remaining inventory as well as remaining time periods are considered. Figure 1(a)-(d) 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 four figures show increasing levels of risk aversion.

Figure 1(e)-(g) display seat and inventory dependent policies. This becomes 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 based on the marginal seat values of the risk-neutral solutions and of the risk-sensitive recursive solutions were analysed between the policies π^{β} and $\pi^{\beta,\nu}$ first. Both use a discount factor in order to relax the decision if an arrival should be accepted or rejected. Figure 2 shows the averaged values of revenue versus standard deviation and of revenue versus CVaR. We used $\alpha = 5\%$ for the CVaR computations in this paper. The policy $\pi^{\beta,\nu}$ slightly outperformed policy π^{β} in both evaluated measures, though the difference was not remarkable at all. This is recognisable as the graph of policy $\pi^{\beta,\nu}$ embeds (is above) the graph of policy π^{β}. Generally, both policies could achieve nearly the same results but responded with different sensitivity when changing the parameter β. In order to achieve a similar average-standard-deviation pair or a similar average-CVaR pair, the parameter β of policy π^{β} had to be greater than of policy $\pi^{\beta,\nu}$. In particular, this was the case with further experiments with a high load factor.

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 $\pi^{\beta_{\kappa_1,\kappa_2},\nu}$ were also negligible. There is little visible difference between the graphs. The results were analogous to those of the comparison of π^{β} and $\pi^{\beta,\nu}$. Similar average-standard-deviation pair or average-CVaR pairs were achieved when the risk sensitivity parameter κ_2 of policy $\pi^{\beta_{\kappa_1,\kappa_2}}$ was greater than of policy$\pi^{\beta_{\kappa_1,\kappa_2},\nu}$. Further experiments with a high load factor showed that, too.

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

We compared 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 used $\pi^{\beta,\nu}$ as representative of both policies which evaluate discounted marginal seat values because it performed at least as well as the other. We used $\pi^{\beta_{\kappa_1,\kappa_2},\nu}$ as a representative of the selling-rate dependent policies $\pi^{\beta_{\kappa_1,\kappa_2}}$ and $\pi^{\beta_{\kappa_1,\kappa_2},\nu}$ as both showed the same performance.

Figure 4(a) shows that the policies $\pi^{\beta,\nu}$ and $\pi^{\beta_{\kappa_1,\kappa_2},\nu}$ had similar volatility. There is no visible difference between $\pi^{\beta_{\kappa_1,\kappa_2},\nu}$ compared to $\pi^{\beta,\nu}$ in terms of standard deviation. The second selling-rate dependent policy $\pi^{\beta,\mathbb{1},\nu}$ that used an indicator function had a higher volatility. This volatility became greater with increasing level of risk aversion (decreasing β). In further experiments, the risk aversion of every policy yielded only a significant effect for greater load factors.

The CVaR measure revealed a different result in terms of downside risk as shown in Figure 4(b). The policy $\pi^{\beta,\nu}$ achieved 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 $\pi^{\beta_{\kappa_1,\kappa_2},\nu}$ performed well here for all level of risk sensitivity. It outperformed policy $\pi^{\beta,\nu}$. On the other hand, policy $\pi^{\beta,\mathbb{1},\nu}$ displayed 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 dropped down significantly. This drop coincided with its behaviour for the standard deviation. However, as a policy which just switches risk aversion on or off, its results were very good for moderate levels of risk aversion. However, this policy had limitations which were apparent when an increase of its risk sensitivity did no longer affect the results.

4.3.3. Comparison of Exponential Utility and Selling-Rate Dependent Policies

Finally, we compared the selling-rate dependent policies with the policy using the exponential utility function. This is demonstrated in Figure 5. It is notable that there was a very similar behaviour between policies $\pi^{\beta_{\kappa_1,\kappa_2},\nu}$ and $\pi^{*\gamma}$. Their results differed slightly only for a high level of risk aversion. This is observable in Figure 5(b) on the lower right hand side. Policy $\pi^{\beta,\mathbb{1},\nu}$

could only follow the other policies up to a certain level of risk sensitivity as already mentioned above. In particular, policy $\pi^{\beta,\mathbb{1},\nu}$ did not performed well for high load factors.

5. Conclusions

Several risk-averse policies for the dynamic capacity control problem were compared with respect to mean revenue versus standard deviation and mean revenue versus conditional-value-at-risk. We have shown that there were only small differences between the risk-sensitive policies employing a discounted decision rule. The results in terms of the proposed risk measures were similar and did not depend on the computation of a complete dynamic programming solution for each level of risk aversion. We had the same situation for the two comparable policies which used a selling-rate dependent policy employing a hyperbolic tangent function. In both cases, there was 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 depending on the current selling-rate. This policy kept up with a previous proposed selling-rate dependent policy and a policy using an exponential utility for a wide range of moderate levels of risk sensitivity in terms of down-side risk. Although we used the selling-rate as a variable for (de)activation of the risk aversion, other variables, e.g., external financial constraints, are feasible for its control.

Finally, we have shown that it is adequate to apply the decision rules on the marginal seat values of the risk-neutral solution to achieve at least similar results to policies that use decision rules based on the marginal seat values of distinct dynamic programming solution of risk-averse policies.

From a practical point of view, the advantages are as follows. First, only the computation of the solution of one dynamic program (the risk-neutral solution) is needed for applying different levels of risk aversions. A practitioner can use this risk-neutral solution to experiment with different levels of risk aversion in order to determine which level is appropriate in terms of chosen measures. Thus, the computational effort of solving further dynamic programs is saved. Second, only the risk-neutral solution needs to be stored in memory even if several levels of risk aversion have to be considered in the same problem setting, e.g., different risk levels for various seasons. The additional storage of the risk levels is sufficient in such cases. The benefits grow with increasing problem size. Third, risk aversion can be easily integrated in the risk-neutral solutions in a readily understandable way. To this end, a small modification of the decision rule of an established risk-neutral implementation is sufficient.

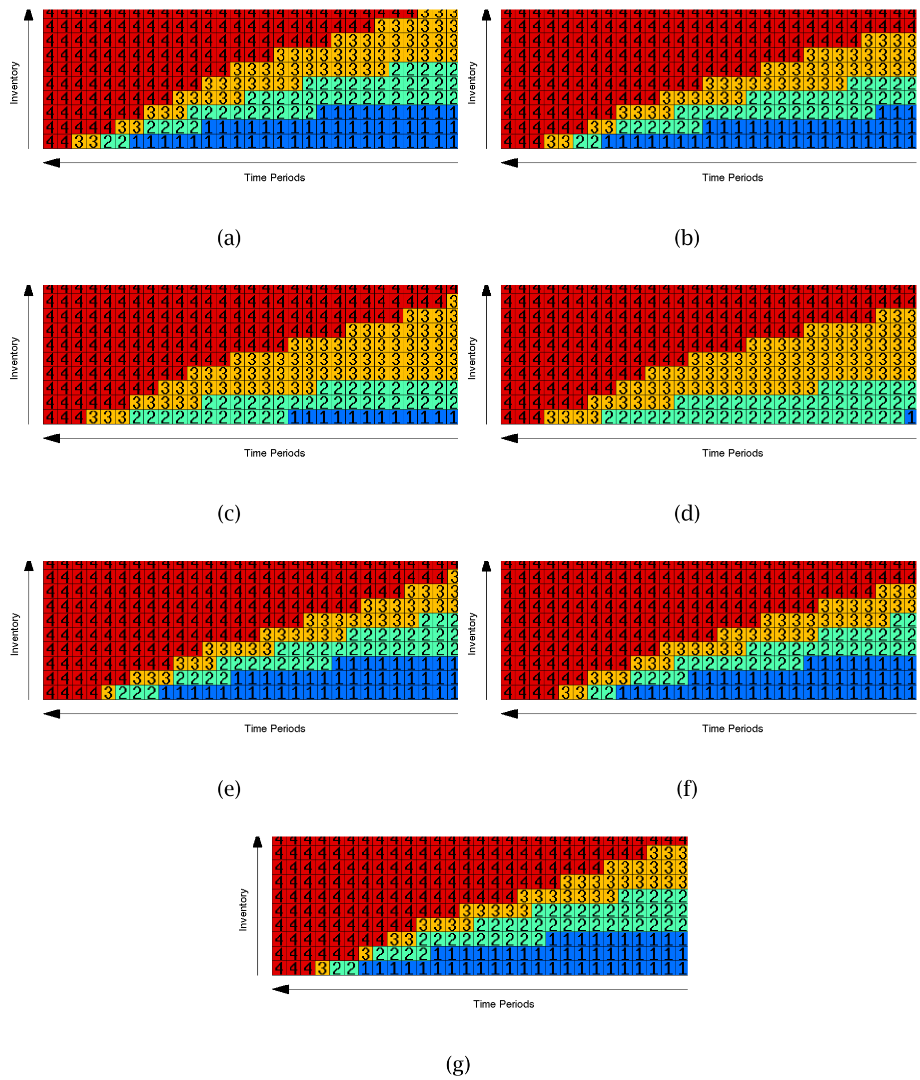

Figure 1: Protection levels generated by the different policies (inventory on x-axis, time periods on y-axis): risk-neutral policy π^* (a), policy using an exponential utility $\pi^{*,\gamma=0.005}$ (b), policy with discounting on risk-neutral solution $\pi^{\beta=0.8,\nu}$ (c), policy with discounting $\pi^{\beta=0.8}$ (d), policy with selling-dependency on risk-neutral solution $\pi^{\beta_{\kappa_1=0.5,\kappa_2=0.8},\nu}$ (e), policy with selling-dependency $\pi^{\beta_{\kappa_1=0.5,\kappa_2=0.8}}$ (f), and policy with selling-rate dependent switch on the risk-neutral solution $\pi^{\beta=0.8,\mathbb{1}}$ (g).

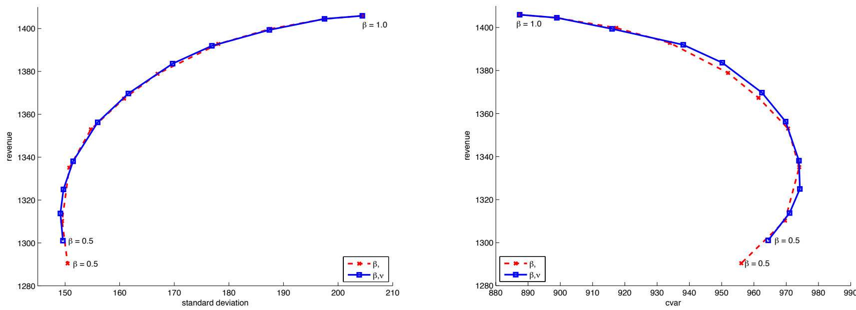

(a) (b)

Figure 2: Comparison of the policies π^{β} and $\pi^{\beta,\nu}$ using revenue vs. standard deviation (a) and revenue vs. conditional-value-of-risk with $\alpha = 5\%$ (b). The value of β goes from 0.5 to 1.0 by a step size of 0.05.

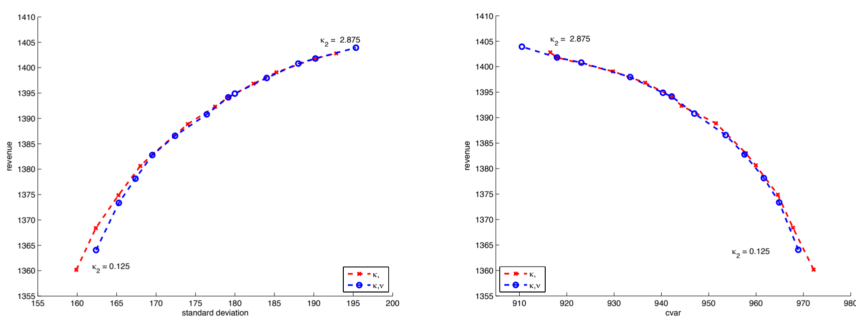

(a) (b)

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 with $\alpha = 5\%$ (b). The value of κ_1 is fixed at 0.5 and κ_2 goes from 0.125 to 2.875 with step size of 0.25.

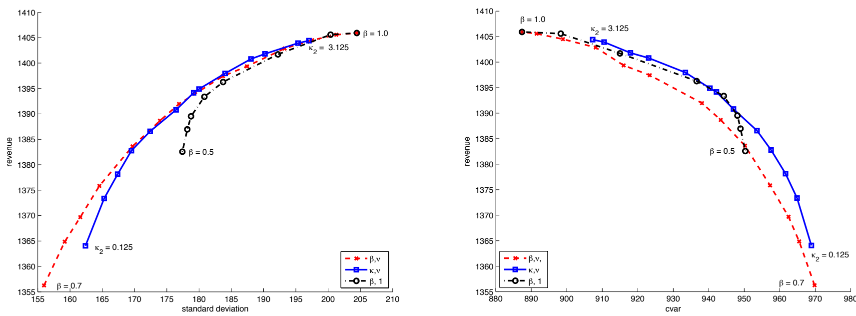

(a) (b)

Figure 4: Comparison of the polices $\pi^{\beta,\nu}$, $\pi^{\beta_{\kappa_1,\kappa_2},\nu}$ and $\pi^{\beta,\mathbb{1},\nu}$ using revenue vs. standard deviation (a) and revenue vs. conditional-value-of-risk with $\alpha = 5\%$. (b). Policy $\pi^{\beta,\nu}$ uses $\beta \in [0.7,\ldots,1.0]$ with step size 0.05. Policy $\pi^{\beta_{\kappa_1,\kappa_2},\nu}$ has κ_1 fixed at 0.5 and κ_2 from 0.125 to 3.125 with step size 0.25. Policy $\pi^{\beta,\mathbb{1},\nu}$ uses $\beta \in [0.5,\ldots,1.0]$ with step size 0.05.

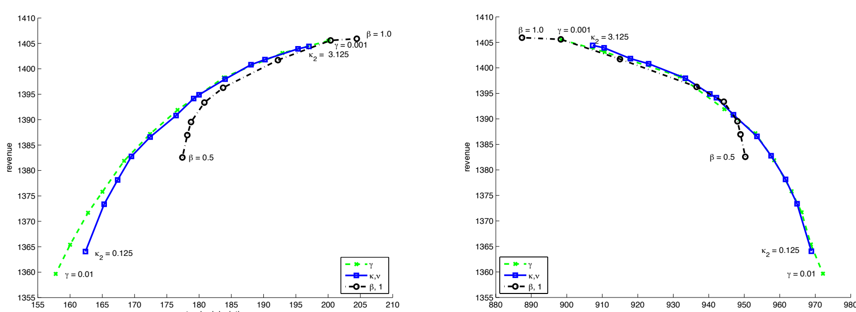

(a) (b)

Figure 5: Comparison of the policies $\pi^{*\gamma}$, $\pi^{\beta_{\kappa_1,\kappa_2},\nu}$ and $\pi^{\beta,\mathbb{1},\nu}$ using revenue vs. standard deviation (a) and revenue vs. conditional-value-of-risk with $\alpha = 5\%$ (b). Policy $\pi^{*\gamma}$ has $\gamma \in [0.001,\ldots,0.01]$ with step size 0.001. Policy $\pi^{\beta_{\kappa_1,\kappa_2},\nu}$ has κ_1 fixed at 0.5 and κ_2 from 0.125 to 3.125 with step size 0.25. Policy $\pi^{\beta,\mathbb{1},\nu}$ uses $\beta \in [0.5,\ldots,1.0]$ with step size 0.05.

References

[1] P. J. H. Schoemaker, The expected utility model: Its variants, purposes, evidence and limitations, Journal of Economic Literature 20 (1982) 529–563.

[2] J.-P. Chavas, Risk Analysis in Theory and Practice, Elsevier Academic Press, 2004.

[3] D. P. Bertsekas, Dynamic Programming and Optimal Control, Vol. I, 3rd ed., Athena Scientific, 2005.

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

[5] T. Lee, M. Hersh, A model for dynamic airline seat inventory controls with multiple seat bookings, Transportation Science 27 (1993) 252–265.

[6] C. Barz, K.-H. Waldmann, Risk-sensitive capacity control in revenue management, Mathematical Methods of Operations Research 65 (2007) 565–579.

[7] K. Huang, K. Chang, A model for airline seat control considering revenue uncertainty and risk, Journal of Revenue and Pricing Management 10 (2011) 161–171.

[8] K. Talluri, G. van Ryzin, The Theory and Practice of Revenue Management, Springer, 2005.

[9] W.-C. Chiang, J. Chen, X. Xu, An overview of research on revenue management: Current issues and future research, International Journal of Revenue Management 1 (2007) 97–128.

[10] Y. Feng, B. Xiao, Maximizing revenues of perishable assets with a risk factor, Operations Research 47 (1999) 337–341.

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

[12] D. Mitra, Q. Wang, Stochastic traffic engineering for demand uncertainty and risk-aware network revenue management, IEEE/ACM Transactions on Networking 13 (2005) 221–233.

[13] M. Koenig, J. Meissner, List pricing versus dynamic pricing: Impact on the revenue risk, European Journal of Operational Research 204 (2010) 505–512.

[14] A. Ben-Tal, A. Nemirovski, Robust convex optimization, Mathematics of Operations Research 23 (1998) 769–805.

[15] G. Perakis, G. Roels, Robust Controls for Network Revenue Management, Technical Report, Decisions, Operations, and Technology Management, University of California, 2007.

[16] A. Thiele, Single-Product Pricing via Robust Optimization, Technical Report, Lehigh University, 2006.

[17] A. Lim, J. Shanthikumar, Relative entropy, exponential utility, and robust dynamic pricing, Operations Research 55 (2007) 198–214.

[18] J. M. Mulvey, R. J. Vanderbei, S. A. Zenios, Robust optimization of large-scale systems, Operations Research 43 (1995) 264–281.

[19] K.-K. Lai, W.-L. Ng, A stochastic approach to hotel revenue optimization, Computers and Operations Research 32 (2005) 1059–1072.

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

[21] P. P. Beloba, Application of an probabilistic decision model to airline seat inventory control, Operations Research 37 (1989) 183–197.

[22] Y. Feng, B. Xiao, A risk-sensitive model for managing perishable products, Operations Research 56 (2008) 1305–1311.

[23] C. Barz, Risk-Averse Capacity Control in Revenue Management, Ph.D. thesis, University of Karlsruhe, 2007.

[24] T. Volling, D. E. Akyol, K. Wittek, T. S. Spengler, A two-stage bid-price control for make-to-order revenue management, Computers and Operations Research 39 (2012) 1021–1032.

[25] R. A. Howard, J. E. Matheson, Risk-sensitive markov decision processes, Management Science 18 (1972) 356–369.

[26] E. Luciano, L. Peccati, D. M. Cifarelli, VaR as a risk measure for multiperiod static inventory models, International Journal of Production Economics 81–82 (2003) 375–284.

[27] S. Ahmed, U. Cakmak, A. Shapiro, Coherent risk measures in inventory problems, European Journal of Operational Research 182 (2006) 226–238.

6. Appendix

Further Numerical Results

We investigated the proposed risk-averse policies by several experiments. We defined 12 groups which differed by their fare structures and request probabilities. The initial capacity was $C = 10$ and the remaining time was $N = 30$ for all experiments. Furthermore, we divided the time periods in five equal parts which each consisted of six periods with fixed request probabilities.

In our experiments, we generated random request probabilities for four classes which were distinguished by their load factor, see below. A class consisted of 100 experiments with different request probabilities. We simulated an experiment by running $1,000$ sample runs with the same request probabilities in Monte Carlo manner.

We computed a load factor ρ with the request probabilities of an experiment as $\rho = \frac{1}{C}\sum_{n=0}^{N}\sum_{j=1}^{k} p_{jn}$. An experiment was assigned to a class by its load factor. Each class contained only experiments with load factors of a certain range: $[0.75, 1.00)$, $[1.00, 1.25)$, $[1.25, 1.50)$ and $[1.50, 1.75)$. We divided each range into 10 smaller intervals of size 0.025 and accepted only 10 experiments for each interval so that the distribution of the load factors became more uniform in a class. Policies to be compared were applied to the same sample run. Then, the results were summarised for each class.

Further, we generated experiments with three different fare structures which we denoted as scenario. Scenario class $S1$ had three fares $200, 150, 100$, scenario class $S2$ had four fares $200, 150, 120, 80$ and scenario class $S3$ had five fares $240, 200, 160, 120, 80$.

Thus, we classified the experiments regarding scenarios on the one hand and load factors on the other hand. Each table shown in the following contains results of certain policies applied to a certain scenario class. The classification regarding the load factors is shown in each table.

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

The results of Table 2 to Table 4 show that there were only small differences between the policies π^{β} and $\pi^{\beta,\nu}$. The results of the three scenario classes were analogous. The effect of using a risk-sensitive policy became noticeable with increasing load factor. The difference between the results of both policies was observable with an increasing load factor and an increasing risk sensitivity given by β, too. However, the same results in pairs of mean and standard deviation (mean and $\text{CVaR}_{5\%}$) could be observed when using a greater parameter β for policy π^{β} than for policy $\pi^{\beta,\nu}$.

Comparison of Policies Using Decision Rules with Selling-Rate Dependent on Marginal Seat Values of Risk-Neutral and Risk-Averse Recursive Solution

We had a similar results of the experiments comparing the policies $\pi^{\beta^{\kappa_1,\kappa_2}}$ and $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$ as with the previous comparison. The results are shown in Table 5 to Table 7. The effect of applying the risk-averse policies became visible with increasing the load factor, and the effect between both policies became visible with increasing load factor and increasing risk aversion controlled by κ_2. We fixed κ_1 at 0.5 in all experiments. The same result pairs of mean and standard deviation (mean and $\text{CVaR}_{5\%}$) were obtainable by using different κ_2 parameters for the policies.

Comparison between all Policies

The results of Table 8 to Table 10 show mean, standard deviation and $\text{CVaR}_{5\%}$ of the policies $\pi^{\beta,\nu}$, $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$, $\pi^{*\gamma}$ and $\pi^{\beta,\mathbb{1},\nu}$, where each table shows one of the three scenarios. As each policy had a different parameter which controlled the level of risk-sensitivity, we chose the parameters so that the mean revenue was nearly the same. In this manner, we were able to compare pairs of mean and standard deviation and pairs of mean and $\text{CVaR}_{5\%}$ respectively. We selected a range of mean and associated values which reflected the 'efficient frontier' of mean vs. $\text{CVaR}_{5\%}$. Thus, we increased the levels of risk aversion as long as the $\text{CVaR}_{5\%}$ increased (and the mean decreased).

We observed that results were similar in respect of their classification by the load factor ranges. The performance of the policies depended on the load factor ranges to a greater extend and on scenarios to a lesser extend.

First, we noticed that a risk-sensitive policy did not affect strongly the results for a low load factor in the range $[0.75, 1.00)$ independent of the scenario. The low load factors represented less requests than available capacity. Every request could be accepted and risk considerations did not play a role. Thus, the performance of a policy was regardless of the level of risk sensitivity, too.

Second, we observed that the policies $\pi^{\beta,\nu}$, $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$ and $\pi^{*\gamma}$ performed similar for the load factor ranges $[1.00, 1.25)$ and $[1.25, 1.50)$ independent of the scenario. While increasing their risk-sensitive control parameters, mean and standard deviation went down and $\text{CVaR}_{5\%}$ went up. Differences between the policies were generally insignificant. The greatest differences were for a high level of risk aversion of the load factors in $[1.25, 1.50)$ in Scenario 3: for a similar mean, policy $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$ had a had 1.7% and 2.0% higher standard deviation than policy $\pi^{*\gamma}$ and policy $\pi^{\beta,\nu}$, respectively; and policy $\pi^{*\gamma}$ had 0.9% and 0.4% higher $\text{CVaR}_{5\%}$ than policy $\pi^{\beta,\nu}$ and policy $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$, respectively. For example, policy $\pi^{*\gamma}$ achieved a 4.7% greater $\text{CVaR}_{5\%}$ at the expense of a 1.0% mean revenue with $\gamma = 0.0048$ for the load factors in $[1.25, 1.50)$ in Scenario 3.

Third, the policy $\pi^{\beta,\mathbb{1},\nu}$ kept up with the other policies only up to a moderate level of risk aversion for the load factor ranges $[1.00, 1.25)$ and $[1.25, 1.50)$. For high levels of risk aversion, we could not adjust the mean revenues of the policy $\pi^{\beta,\mathbb{1},\nu}$ to every of the mean revenues of the other policies by modifying the parameter β. The values of β shown in the tables demonstrate how the policy $\pi^{\beta,\mathbb{1},\nu}$ came to a saturation. However, we achieved, e.g., 4.0% greater $\text{CVaR}_{5\%}$ at the expense of 1.1% revenue using policy $\pi^{\beta,\mathbb{1},\nu}$ with $\beta = 0.7$ for load factors in $[1.25, 1.50)$ in Scenario 3.

Fourth, the results of the policies for the load factors range $[1.50, 1.75)$ showed the boundaries of the useful ranges of levels of risk aversion. The increase of the level of risk aversion reduced the mean revenue and the standard deviation for all policies with the exception of policy $\pi^{\beta,\mathbb{1},\nu}$. However, the increase of the $\text{CVaR}_{5\%}$ did not work in accordance with raising the level of risk sensitivity in Scenarios 2 and 3. The $\text{CVaR}_{5\%}$ increased up to a certain level and decreased then although the level of risk aversion was incremented. The differences between the policies $\pi^{\beta,\nu}$, $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$ and $\pi^{*\gamma}$ were negligible. The best achieved results were a decrease of the standard deviation by 20% accompanied by a decrease of the mean by 3.5% (policy $\pi^{\beta,\nu}$, $\beta = 0.7$, Scenario 3), and an increase of the $\text{CVaR}_{5\%}$ by 4.9% accompanied by a decrease of the mean by 1.9% (policy $\pi^{*\gamma}$, $\gamma = 0.008$, Scenario 3). The policy $\pi^{\beta,\mathbb{1},\nu}$ performed comparably with the other policies up to a moderate level of risk sensitivity. For comparison, e.g., policy $\pi^{\beta,\mathbb{1},\nu}$ achieved a reduction of the standard deviation by 7.4% and an increment of the $\text{CVaR}_{5\%}$ by 3.1% coincided with a decrease of the mean by 0.5% ($\beta = 0.85$, Scenario 3).

Exemplary Comparison of Risk-averse Policies' Revenue Distribution

In order to visualise the different behaviour of the risk-averse policies, we provide revenue distributions of an example setting in Figure 6. We chose an experiment from Scenario class S3 with a load factor of 1.23 and adjusted the risk levels of the policies so that their mean revenues are similar as shown in Table 11. We observe that the distributions between the policies π^{β} and $\pi^{\beta,\nu}$ are approximate equal and between the policies $\pi^{\beta^{\kappa_1,\kappa_2}}$ and $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$ are equal for the adjusted risk levels. The different $\text{CVaR}_{5\%}$s as given in Table 11 are observable by looking at the lower end of the distributions, policies π^{β} and $\pi^{\beta,\nu}$ produced apparently more lower revenues here. The higher standard deviation of policy $\pi^{\beta,\mathbb{1},\nu}$ compared to the others is discernible, too.

β	Policy π^{β} mean	std. dev.	$CVaR_{5\%}$	Policy $\pi^{\beta,\nu}$ mean	std. dev.	$CVaR_{5\%}$
Scenario $S1$, $\rho \in [0.75, 1.00)$, mean(V^*)=1247.02						
1.0	1247.30	284.09	612.72	1247.30	284.09	612.72
0.9	1246.82	282.75	612.74	1246.87	282.80	612.74
0.8	1245.55	281.23	612.75	1245.78	281.46	612.75
0.7	1243.77	279.68	612.75	1244.34	280.12	612.75
0.6	1242.43	278.74	612.75	1242.88	279.01	612.75
0.5	1242.42	278.73	612.75	1242.42	278.73	612.75
Scenario $S1$, $\rho \in [1.00, 1.25)$, mean(V^*)=1461.07						
1.0	1460.79	207.99	937.54	1460.79	207.99	937.54
0.9	1457.39	198.67	946.87	1457.83	199.16	946.65
0.8	1449.29	191.61	950.74	1451.30	192.86	950.67
0.7	1438.75	186.67	952.10	1442.86	188.37	952.11
0.6	1430.13	184.46	952.25	1433.80	185.37	952.42
0.5	1428.99	184.48	951.92	1428.99	184.48	951.92
Scenario $S1$, $\rho \in [1.25, 1.50)$, mean(V^*)=1593.57						
1.0	1593.75	170.70	1163.77	1593.75	170.70	1163.77
0.9	1583.60	151.54	1202.21	1586.27	154.04	1199.04
0.8	1559.59	137.74	1224.76	1568.76	141.49	1222.62
0.7	1523.82	131.83	1221.00	1543.18	134.41	1227.90
0.6	1494.17	132.75	1200.06	1511.85	132.33	1213.60
0.5	1492.04	133.30	1197.21	1492.04	133.30	1197.21
Scenario $S1$, $\rho \in [1.50, 1.75)$, mean(V^*)=1687.37						
1.0	1687.08	146.79	1308.29	1687.08	146.79	1308.29
0.9	1673.52	128.78	1345.96	1678.61	131.88	1343.45
0.8	1634.69	117.25	1360.76	1656.63	122.28	1360.65
0.7	1569.62	116.14	1324.11	1615.23	116.93	1354.88
0.6	1506.92	123.63	1255.95	1550.36	119.28	1301.76
0.5	1500.16	125.14	1246.34	1500.16	125.14	1246.34

Table 2: Comparison of policies depending on marginal seat values of risk-neutral and risk-averse recursive solution: Scenario 1

β	Policy π^{β} mean	std. dev.	$CVaR_{5\%}$	Policy $\pi^{\beta,\nu}$ mean	std. dev.	$CVaR_{5\%}$
Scenario $S2$, $\rho \in [0.75, 1.00)$, mean(V^*)=1136.90						
1.0	1136.33	270.83	537.50	1136.33	270.83	537.50
0.9	1135.81	269.33	537.60	1135.81	269.32	537.60
0.8	1134.71	267.71	537.65	1134.82	267.81	537.65
0.7	1133.10	266.09	537.65	1133.37	266.34	537.65
0.6	1131.13	264.57	537.65	1131.82	265.09	537.65
0.5	1130.37	264.09	537.65	1130.60	264.21	537.65
Scenario $S2$, $\rho \in [1.00, 1.25)$, mean(V^*)=1339.89						
1.0	1339.12	214.23	817.09	1339.12	214.23	817.09
0.9	1336.55	205.55	827.20	1336.63	205.65	827.12
0.8	1329.76	198.32	832.66	1330.42	198.71	833.03
0.7	1320.47	192.77	835.94	1321.61	193.33	835.91
0.6	1308.79	188.72	836.05	1312.33	189.83	836.19
0.5	1302.86	187.61	835.45	1305.49	188.00	835.88
Scenario $S2$, $\rho \in [1.25, 1.50)$, mean(V^*)=1480.69						
1.0	1480.17	180.52	1031.57	1480.17	180.52	1031.57
0.9	1474.55	167.41	1055.79	1475.16	167.96	1055.82
0.8	1457.91	154.67	1078.33	1461.63	156.57	1077.48
0.7	1428.25	145.48	1089.32	1436.05	147.30	1090.43
0.6	1388.16	142.37	1073.38	1402.16	143.22	1081.45
0.5	1364.12	143.83	1052.67	1375.89	143.08	1063.33
Scenario $S2$, $\rho \in [1.50, 1.75)$, mean(V^*)=1571.84						
1.0	1571.55	154.94	1187.65	1571.55	154.94	1187.65
0.9	1561.84	138.24	1220.65	1563.54	139.40	1220.35
0.8	1536.55	128.09	1233.81	1544.62	130.29	1234.12
0.7	1494.76	124.44	1222.82	1516.49	126.16	1232.72
0.6	1428.61	128.50	1164.45	1456.58	127.15	1191.23
0.5	1381.91	134.78	1108.23	1408.37	131.67	1139.30

Table 3: Comparison of policies depending on marginal seat values of risk-neutral and risk-averse recursive solution: Scenario 2

β	mean	Policy π^{β} std. dev.	$CVaR_{5\%}$	mean	Policy $\pi^{\beta,\nu}$ std. dev.	$CVaR_{5\%}$
Scenario $S3$, $\rho \in [0.75, 1.00)$, mean(V^*)=1319.00						
1.0	1319.13	328.83	603.76	1319.13	328.83	603.76
0.9	1318.66	326.85	604.04	1318.65	326.80	604.06
0.8	1317.27	324.64	604.23	1317.31	324.68	604.23
0.7	1315.28	322.48	604.27	1315.34	322.54	604.27
0.6	1312.79	320.33	604.27	1313.17	320.65	604.28
0.5	1310.28	318.57	604.29	1311.09	319.12	604.29
Scenario $S3$, $\rho \in [1.00, 1.25)$, mean(V^*)=1578.26						
1.0	1577.45	272.48	925.54	1577.45	272.48	925.54
0.9	1575.10	263.11	936.31	1575.11	263.00	936.78
0.8	1567.78	253.41	945.63	1567.90	253.37	946.10
0.7	1557.01	245.01	951.25	1557.19	245.01	952.37
0.6	1543.22	238.67	954.14	1544.26	239.11	954.62
0.5	1527.66	234.15	953.81	1532.45	235.37	954.50
Scenario $S3$, $\rho \in [1.25, 1.50)$, mean(V^*)=1733.60						
1.0	1733.78	230.66	1168.00	1733.78	230.66	1168.00
0.9	1729.04	216.60	1192.18	1729.28	216.84	1192.36
0.8	1713.92	203.04	1213.17	1716.30	204.51	1213.12
0.7	1686.46	191.08	1229.07	1690.58	192.84	1230.58
0.6	1647.76	184.33	1229.68	1650.97	185.39	1231.06
0.5	1594.78	183.46	1199.38	1613.39	183.81	1212.51
Scenario $S3$, $\rho \in [1.50, 1.75)$, mean(V^*)=1863.28						
1.0	1863.91	202.06	1358.43	1863.91	202.06	1358.43
0.9	1854.23	181.54	1399.02	1855.78	183.05	1397.49
0.8	1827.97	166.25	1422.30	1833.99	168.69	1421.39
0.7	1785.68	158.89	1420.61	1798.76	160.90	1424.92
0.6	1729.68	160.38	1383.46	1746.32	160.70	1394.94
0.5	1636.52	169.31	1287.41	1675.66	166.16	1327.84

Table 4: Comparison of policies depending on marginal seat values of risk-neutral and risk-averse recursive solution: Scenario 3

κ_2	Policy $\pi^{\beta^{\kappa_1,\kappa_2}}$ mean	std. dev.	$CVaR_{5\%}$	Policy $\pi^{\beta^{\kappa_1,\kappa_2,\nu}}$ mean	std. dev.	$CVaR_{5\%}$
Scenario $S1$, $\rho \in [0.75, 1.00)$, mean(V^*)=1247.02						
8.0	1247.30	284.08	612.72	1247.30	284.08	612.72
4.0	1247.27	283.96	612.72	1247.27	283.96	612.72
2.0	1247.09	283.35	612.74	1247.10	283.36	612.74
1.0	1246.58	282.56	612.74	1246.60	282.58	612.74
0.5	1245.99	281.86	612.74	1246.07	281.94	612.74
Scenario $S1$, $\rho \in [1.00, 1.25)$, mean(V^*)=1461.07						
8.0	1460.78	207.97	937.59	1460.78	207.97	937.59
4.0	1460.63	206.44	939.65	1460.63	206.45	939.65
2.0	1458.57	200.48	946.46	1458.75	200.75	946.21
1.0	1454.13	195.57	949.98	1454.78	196.16	949.74
0.5	1450.44	193.02	951.12	1451.47	193.77	950.95
Scenario $S1$, $\rho \in [1.25, 1.50)$, mean(V^*)=1593.57						
8.0	1593.76	170.50	1164.32	1593.76	170.50	1164.32
4.0	1592.85	165.09	1178.36	1592.93	165.29	1178.01
2.0	1581.35	150.39	1212.34	1583.36	152.16	1208.95
1.0	1564.71	142.50	1224.79	1569.59	144.68	1222.63
0.5	1553.61	139.38	1226.14	1560.56	141.90	1225.56
Scenario $S1$, $\rho \in [1.50, 1.75)$, mean(V^*)=1687.37						
8.0	1687.07	146.63	1308.97	1687.07	146.63	1308.97
4.0	1684.93	140.45	1329.29	1685.42	140.87	1328.34
2.0	1661.69	127.94	1357.65	1669.66	130.65	1355.93
1.0	1631.60	123.50	1353.94	1645.43	126.32	1356.71
0.5	1612.33	122.61	1343.58	1628.99	125.27	1349.43

Table 5: Comparison of policies with decisions using selling-rate depending on marginal seat values of risk-neutral and risk-averse recursive solution: Scenario 1

κ_2	Policy $\pi^{\beta^{\kappa_1,\kappa_2}}$ mean	std. dev.	$CVaR_{5\%}$	Policy $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$ mean	std. dev.	$CVaR_{5\%}$
Scenario $S2$, $\rho \in [0.75, 1.00)$, mean(V^*)=1136.90						
8.0	1136.33	270.83	537.50	1136.33	270.83	537.50
4.0	1136.33	270.65	537.53	1136.33	270.65	537.53
2.0	1136.08	269.94	537.56	1136.08	269.95	537.56
1.0	1135.60	268.96	537.64	1135.62	269.00	537.64
0.5	1134.99	268.16	537.65	1135.05	268.23	537.65
Scenario $S2$, $\rho \in [1.00, 1.25)$, mean(V^*)=1339.89						
8.0	1339.13	214.19	817.16	1339.13	214.19	817.16
4.0	1339.05	212.42	819.80	1339.05	212.42	819.78
2.0	1336.68	206.04	828.50	1336.78	206.22	828.34
1.0	1331.88	200.64	833.40	1332.54	201.19	832.92
0.5	1327.96	197.91	834.69	1329.05	198.69	834.30
Scenario $S2$, $\rho \in [1.25, 1.50)$, mean(V^*)=1480.69						
8.0	1480.14	180.47	1031.74	1480.14	180.47	1031.74
4.0	1479.56	176.46	1043.79	1479.57	176.53	1043.53
2.0	1469.10	163.83	1074.47	1471.28	165.42	1071.33
1.0	1451.58	155.08	1087.28	1456.06	157.14	1085.43
0.5	1439.21	151.99	1088.64	1445.02	154.06	1088.01
Scenario $S2$, $\rho \in [1.50, 1.75)$, mean(V^*)=1571.84						
8.0	1571.53	154.86	1187.88	1571.53	154.86	1187.88
4.0	1570.24	149.92	1204.56	1570.50	150.17	1203.89
2.0	1553.90	139.43	1229.45	1558.50	140.89	1228.69
1.0	1524.31	134.61	1225.58	1535.24	136.87	1228.07
0.5	1505.45	133.43	1215.89	1518.38	135.79	1220.73

Table 6: Comparison of policies with decisions using selling-rate depending on marginal seat values of risk-neutral and risk-averse recursive solution: Scenario 2

κ_2	Policy $\pi^{\beta^{\kappa_1,\kappa_2}}$ mean	std. dev.	$CVaR_{5\%}$	Policy $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$ mean	std. dev.	$CVaR_{5\%}$
Scenario $S3$, $\rho \in [0.75, 1.00)$, mean(V^*)=1319.00						
8.0	1319.13	328.83	603.76	1319.13	328.83	603.76
4.0	1319.15	328.63	603.79	1319.15	328.63	603.79
2.0	1318.86	327.53	603.96	1318.86	327.53	603.96
1.0	1318.12	326.01	604.07	1318.16	326.09	604.07
0.5	1317.35	324.99	604.14	1317.47	325.13	604.14
Scenario $S3$, $\rho \in [1.00, 1.25)$, mean(V^*)=1578.26						
8.0	1577.44	272.37	925.79	1577.44	272.37	925.79
4.0	1577.23	270.16	929.51	1577.23	270.16	929.51
2.0	1574.36	261.78	940.96	1574.47	262.01	940.79
1.0	1567.84	254.07	948.57	1568.39	254.65	948.26
0.5	1562.70	249.88	951.23	1563.73	250.85	950.75
Scenario $S3$, $\rho \in [1.25, 1.50)$, mean(V^*)=1733.60						
8.0	1733.76	230.52	1168.36	1733.76	230.52	1168.36
4.0	1733.07	225.93	1180.62	1733.10	226.04	1180.31
2.0	1722.49	212.04	1212.66	1724.38	213.52	1210.24
1.0	1701.15	201.41	1229.16	1706.02	203.61	1227.25
0.5	1686.68	196.75	1232.66	1693.36	199.55	1231.49
Scenario $S3$, $\rho \in [1.50, 1.75)$, mean(V^*)=1863.28						
8.0	1863.91	201.92	1358.94	1863.91	201.92	1358.94
4.0	1862.29	195.44	1378.36	1862.46	195.93	1377.02
2.0	1842.83	180.24	1413.19	1847.55	182.57	1411.08
1.0	1807.31	173.68	1410.87	1820.13	176.82	1413.72
0.5	1781.82	172.53	1397.68	1798.00	175.13	1403.50

Table 7: Comparison of policies with decisions using selling-rate depending on marginal seat values of risk-neutral and risk-averse recursive solution: Scenario 3

Scenario		mean	std. dev.	$CVaR_{5\%}$		mean	std. dev.	$CVaR_{5\%}$
Scenario $S1$, $\rho \in [0.75, 1.00)$	β		Policy $\pi^{\beta,\nu}$		κ_2		Policy $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$	
	1.000	1247.30	284.09	612.72	8.000	1247.30	284.08	612.72
	0.950	1247.14	283.39	612.74	4.000	1247.27	283.96	612.72
	0.900	1246.87	282.80	612.74	2.000	1247.10	283.36	612.74
	0.850	1246.36	282.12	612.74	1.500	1246.95	283.05	612.74
	0.800	1245.78	281.46	612.75	1.000	1246.60	282.58	612.74
	γ		Policy $\pi^{*\gamma}$		β		Policy $\pi^{\beta,\mathbb{1},\nu}$	
	0.0001	1247.31	284.04	612.72	1.000	1247.30	284.09	612.72
	0.0010	1247.21	283.64	612.74	0.900	1247.27	284.05	612.72
	0.0030	1246.95	282.91	612.74	0.700	1247.20	283.98	612.72
	0.0040	1246.74	282.62	612.74	0.600	1247.20	283.98	612.72
	0.0080	1245.90	281.62	612.74	0.500	1247.20	283.98	612.72
Scenario $S1$, $\rho \in [1.00, 1.25)$	β		Policy $\pi^{\beta,\nu}$		κ_2		Policy $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$	
	1.000	1460.79	207.99	937.54	8.000	1460.78	207.97	937.59
	0.950	1459.80	203.05	943.18	3.000	1460.19	204.28	942.64
	0.920	1458.72	200.64	945.51	2.000	1458.75	200.75	946.21
	0.890	1457.23	198.36	947.21	1.500	1457.04	198.55	948.09
	0.850	1454.86	195.85	949.05	1.000	1454.78	196.16	949.74
	γ		Policy $\pi^{*\gamma}$		β		Policy $\pi^{\beta,\mathbb{1},\nu}$	
	0.0001	1460.80	207.71	938.04	1.000	1460.79	207.99	937.54
	0.0020	1459.61	202.20	944.58	0.950	1460.42	205.90	941.06
	0.0030	1458.62	200.11	947.06	0.800	1459.51	204.16	943.37
	0.0040	1457.21	197.99	948.87	0.700	1459.23	204.04	943.37
	0.0060	1454.65	195.42	950.50	0.500	1459.13	204.00	943.35
Scenario $S1$, $\rho \in [1.25, 1.50)$	β		Policy $\pi^{\beta,\nu}$		κ_2		Policy $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$	
	1.000	1593.75	170.70	1163.77	8.000	1593.76	170.50	1164.32
	0.950	1591.59	161.87	1182.67	4.500	1593.39	167.32	1173.02
	0.900	1586.27	154.04	1199.04	2.500	1587.43	156.05	1200.04
	0.875	1582.42	150.30	1207.14	2.000	1583.36	152.16	1208.95
	0.825	1573.98	144.25	1218.02	1.500	1577.66	148.14	1217.23
	γ		Policy $\pi^{*\gamma}$		β		Policy $\pi^{\beta,\mathbb{1},\nu}$	
	0.001	1593.37	166.33	1174.55	1.000	1593.75	170.70	1163.77
	0.002	1592.09	162.13	1185.38	0.950	1592.44	164.27	1180.38
	0.004	1586.81	153.86	1205.71	0.850	1586.13	156.04	1201.75
	0.005	1583.60	150.77	1212.33	0.800	1583.02	154.16	1207.13
	0.007	1576.50	145.56	1222.59	0.500	1578.11	153.71	1208.02
Scenario $S1$, $\rho \in [1.50, 1.75)$	β		Policy $\pi^{\beta,\nu}$		κ_2		Policy $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$	
	1.000	1687.08	146.79	1308.29	8.000	1687.07	146.63	1308.97
	0.950	1684.41	138.22	1328.55	5.000	1686.70	143.95	1318.78
	0.900	1678.61	131.88	1343.45	3.000	1681.25	136.51	1341.55
	0.850	1669.69	126.75	1353.94	2.000	1669.66	130.65	1355.93
	0.775	1648.12	120.25	1363.03	1.000	1645.43	126.32	1356.71
	γ		Policy $\pi^{*\gamma}$		β		Policy $\pi^{\beta,\mathbb{1},\nu}$	
	0.0001	1687.07	146.40	1309.72	1.000	1687.08	146.79	1308.29
	0.0030	1684.86	138.38	1334.35	0.950	1685.80	142.06	1324.31
	0.0060	1677.53	131.71	1354.20	0.850	1677.82	135.75	1345.26
	0.0080	1668.97	128.09	1362.04	0.800	1671.52	133.98	1350.77
	0.0120	1645.53	122.58	1364.19	0.500	1643.89	135.71	1337.44

Table 8: Comparison of four risk-sensitive policies: Scenario 1

		mean	std. dev.	$CVaR_{5\%}$		mean	std. dev.	$CVaR_{5\%}$
Scenario $S2$, $\rho \in [0.75, 1.00)$	β		Policy $\pi^{\beta,\nu}$		κ_2		Policy $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$	
	1.000	1136.33	270.83	537.50	8.000	1136.33	270.83	537.50
	0.950	1136.18	270.06	537.56	4.000	1136.33	270.65	537.53
	0.900	1135.81	269.32	537.60	3.000	1136.29	270.44	537.53
	0.850	1135.42	268.60	537.64	2.000	1136.08	269.95	537.56
	0.800	1134.82	267.81	537.65	1.000	1135.62	269.00	537.64
	γ		Policy $\pi^{*\gamma}$		β		Policy $\pi^{\beta,\mathbb{1},\nu}$	
	0.0001	1136.33	270.74	537.53	1.000	1136.33	270.83	537.50
	0.0020	1136.07	269.65	537.63	0.900	1136.28	270.70	537.50
	0.0030	1135.81	269.19	537.64	0.800	1136.25	270.64	537.52
	0.0050	1135.30	268.39	537.65	0.700	1136.24	270.62	537.52
	0.0080	1134.51	267.46	537.65	0.500	1136.24	270.62	537.52
Scenario $S2$, $\rho \in [1.00, 1.25)$	β		Policy $\pi^{\beta,\nu}$		κ_2		Policy $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$	
	1.000	1339.12	214.23	817.09	8.000	1339.13	214.19	817.16
	0.950	1338.47	209.66	823.10	4.000	1339.05	212.42	819.78
	0.900	1336.63	205.65	827.12	2.000	1336.78	206.22	828.34
	0.800	1330.42	198.71	833.03	1.000	1332.54	201.19	832.92
	0.700	1321.61	193.33	835.91	0.100	1325.68	196.61	835.25
	γ		Policy $\pi^{*\gamma}$		β		Policy $\pi^{\beta,\mathbb{1},\nu}$	
	0.0001	1339.11	213.91	817.57	1.000	1339.12	214.23	817.09
	0.0010	1338.73	210.50	822.59	0.950	1338.67	211.51	821.98
	0.0040	1334.51	202.46	832.07	0.750	1336.60	208.24	827.17
	0.0080	1328.50	197.34	835.37	0.600	1336.10	208.06	827.16
	0.0140	1320.96	193.39	836.38	0.500	1336.06	208.05	827.18
Scenario $S2$, $\rho \in [1.25, 1.50)$	β		Policy $\pi^{\beta,\nu}$		κ_2		Policy $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$	
	1.000	1480.17	180.52	1031.57	8.000	1480.14	180.47	1031.74
	0.950	1478.86	174.08	1044.28	4.000	1479.57	176.53	1043.53
	0.900	1475.16	167.96	1055.82	2.500	1475.15	169.01	1062.84
	0.780	1457.65	154.58	1080.64	1.000	1456.06	157.14	1085.43
	0.700	1436.05	147.30	1090.43	0.250	1439.43	152.76	1088.08
	γ		Policy $\pi^{*\gamma}$		β		Policy $\pi^{\beta,\mathbb{1},\nu}$	
	0.000	1480.15	180.19	1032.48	1.000	1480.17	180.52	1031.57
	0.002	1478.51	173.03	1050.85	0.900	1477.68	172.92	1053.08
	0.003	1476.49	169.24	1060.08	0.850	1474.70	169.51	1063.54
	0.008	1455.83	154.66	1088.70	0.600	1457.05	163.66	1076.43
	0.013	1435.27	148.62	1091.91	0.500	1456.64	163.73	1076.08
Scenario $S2$, $\rho \in [1.50, 1.75)$	β		Policy $\pi^{\beta,\nu}$		κ_2		Policy $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$	
	1.000	1571.55	154.94	1187.65	8.000	1571.53	154.86	1187.88
	0.950	1569.05	146.01	1206.70	4.000	1570.50	150.17	1203.89
	0.900	1563.54	139.40	1220.35	3.000	1567.63	146.01	1215.88
	0.800	1544.62	130.29	1234.12	1.500	1549.56	138.82	1230.50
	0.700	1516.49	126.16	1232.72	0.500	1518.38	135.79	1220.73
	γ		Policy $\pi^{*\gamma}$		β		Policy $\pi^{\beta,\mathbb{1},\nu}$	
	0.0001	1571.57	154.62	1188.44	1.000	1571.55	154.94	1187.65
	0.0040	1567.57	143.40	1219.30	0.900	1568.64	148.23	1209.09
	0.0050	1565.49	141.30	1224.70	0.850	1566.59	147.08	1214.41
	0.0100	1546.03	134.10	1237.46	0.700	1549.29	145.78	1218.15
	0.0150	1517.10	130.41	1229.56	0.500	1526.93	149.13	1197.95

Table 9: Comparison of four risk-sensitive policies: Scenario 2

		mean	std. dev.	$CVaR_{5\%}$		mean	std. dev.	$CVaR_{5\%}$
Scenario $S3$, $\rho \in [0.75, 1.00)$	β		Policy $\pi^{\beta,\nu}$		κ_2		Policy $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$	
	1.000	1319.13	328.83	603.76	8.000	1319.13	328.83	603.76
	0.950	1318.95	327.73	603.97	2.000	1318.86	327.53	603.96
	0.900	1318.65	326.80	604.06	1.000	1318.16	326.09	604.07
	0.850	1318.10	325.78	604.14	0.500	1317.47	325.13	604.14
	0.800	1317.31	324.68	604.23	0.250	1316.97	324.54	604.18
	γ		Policy $\pi^{*\gamma}$		β		Policy $\pi^{\beta,\mathbb{1},\nu}$	
	0.0001	1319.13	328.66	603.79	1.000	1319.13	328.83	603.76
	0.0010	1318.93	327.44	604.01	0.800	1318.95	328.39	603.78
	0.0020	1318.58	326.52	604.13	0.700	1318.90	328.34	603.78
	0.0030	1318.12	325.70	604.23	0.600	1318.88	328.31	603.78
	0.0040	1317.57	324.94	604.25	0.500	1318.88	328.31	603.78
Scenario $S3$, $\rho \in [1.00, 1.25)$	β		Policy $\pi^{\beta,\nu}$		κ_2		Policy $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$	
	1.000	1577.45	272.48	925.54	8.000	1577.44	272.37	925.79
	0.950	1576.96	267.92	931.17	3.000	1576.81	267.43	933.63
	0.850	1572.11	258.16	941.87	2.000	1574.47	262.01	940.79
	0.750	1562.90	248.85	949.78	1.000	1568.39	254.65	948.26
	0.680	1554.60	243.56	953.12	0.200	1560.28	248.50	952.18
	γ		Policy $\pi^{*\gamma}$		β		Policy $\pi^{\beta,\mathbb{1},\nu}$	
	0.0001	1577.44	271.96	926.40	1.000	1577.45	272.48	925.54
	0.0010	1576.63	266.39	934.58	0.900	1576.35	266.82	935.19
	0.0020	1574.42	261.13	941.54	0.800	1573.83	262.77	941.59
	0.0050	1564.74	250.33	951.58	0.600	1571.19	260.86	944.08
	0.0090	1554.12	243.68	954.70	0.500	1571.03	260.84	944.07
Scenario $S3$, $\rho \in [1.25, 1.50)$	β		Policy $\pi^{\beta,\nu}$		κ_2		Policy $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$	
	1.000	1733.78	230.66	1168.00	8.000	1733.76	230.52	1168.36
	0.950	1732.55	223.57	1180.54	4.000	1733.10	226.04	1180.31
	0.900	1729.28	216.84	1192.36	3.000	1731.27	221.34	1192.76
	0.800	1716.30	204.51	1213.12	1.500	1716.86	208.67	1219.61
	0.700	1690.58	192.84	1230.58	0.500	1693.36	199.55	1231.49
	γ		Policy $\pi^{*\gamma}$		β		Policy $\pi^{\beta,\mathbb{1},\nu}$	
	0.0001	1733.75	229.89	1169.56	1.000	1733.78	230.66	1168.00
	0.0010	1732.99	224.67	1181.75	0.900	1731.88	223.22	1187.14
	0.0020	1730.87	218.88	1194.56	0.850	1729.70	220.31	1194.98
	0.0048	1716.19	205.04	1224.07	0.700	1714.49	212.81	1214.64
	0.0080	1691.68	194.88	1236.51	0.500	1700.61	211.60	1213.56
Scenario $S3$, $\rho \in [1.50, 1.75)$	β		Policy $\pi^{\beta,\nu}$		κ_2		Policy $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$	
	1.000	1863.91	202.06	1358.43	8.000	1863.91	201.92	1358.94
	0.950	1861.55	192.27	1378.62	4.000	1862.46	195.93	1377.02
	0.900	1855.78	183.05	1397.49	3.000	1858.87	190.06	1393.28
	0.800	1833.99	168.69	1421.39	1.500	1836.40	179.18	1415.38
	0.700	1798.76	160.90	1424.92	0.500	1798.00	175.13	1403.50
	γ		Policy $\pi^{*\gamma}$		β		Policy $\pi^{\beta,\mathbb{1},\nu}$	
	0.0001	1863.89	201.65	1359.42	1.000	1863.91	202.06	1358.43
	0.0020	1861.08	190.81	1386.86	0.900	1859.32	190.13	1391.31
	0.0040	1854.17	181.67	1409.01	0.850	1855.41	186.95	1400.29
	0.0080	1827.65	170.82	1426.01	0.620	1823.45	187.13	1396.07
	0.0110	1798.89	167.00	1419.71	0.500	1799.42	192.15	1369.67

Table 10: Comparison of four risk-sensitive policies: Scenario 3

Policy	mean	std. dev	$CVaR_{5\%}$
$\pi^{\beta=0.81}$	1766.68	238.45	1164.80
$\pi^{\beta=0.8,\nu}$	1766.16	237.78	1165.60
$\pi^{\beta^{\kappa_1=0.5,\kappa_2=1.35}}$	1766.36	237.22	1185.60
$\pi^{\beta^{\kappa_1=0.5,\kappa_2=1.25},\nu}$	1766.36	237.22	1185.60
$\pi^{\beta=0.7,\mathbb{1},\nu}$	1766.20	243.22	1180.00
$\pi^{*\gamma=0.0035}$	1766.16	235.65	1183.20

Table 11: Parameters and values of example of Figure 6.

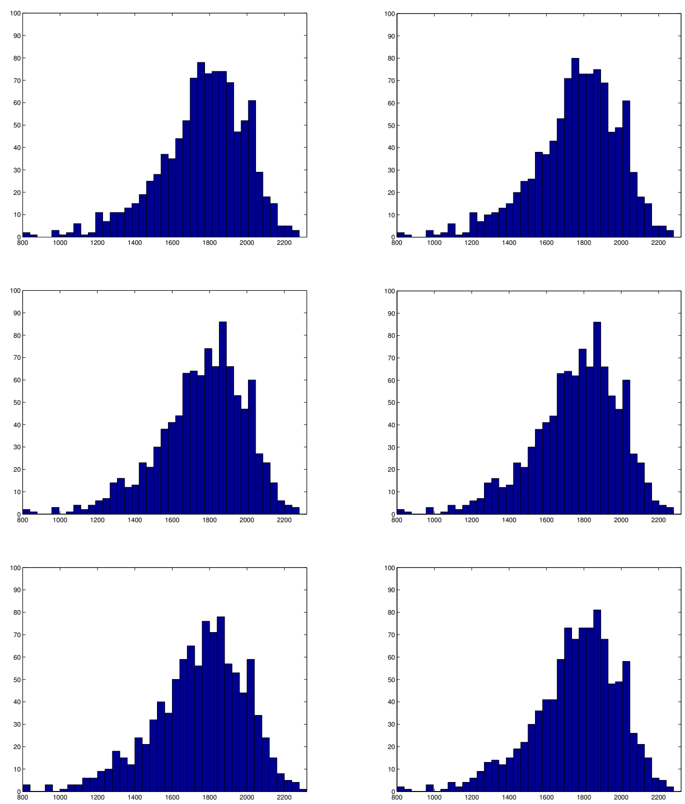

Figure 6: Example: Histrograms of the distributions of revenues achieved by the policies π^{β}, $\pi^{\beta,\nu}$, $\pi^{\beta^{\kappa_1,\kappa_2}}$, $\pi^{\beta^{\kappa_1,\kappa_2},\nu}$, $\pi^{\beta,\mathbb{1},\nu}$, and $\pi^{*\gamma}$ (from top left to bottom right) with risk levels adjusted for similar means.