

Value-At-Risk Optimal Policies for Revenue Management Problems

Matthias Koenig

*Department of Management Science, Lancaster University Management School, United Kingdom.
m.koenig@lancaster.ac.uk*

Joern Meissner

Kuehne Logistics University, Hamburg, Germany. joe@meiss.com

June 12, 2011

Abstract

Consider a single-leg dynamic revenue management problem with fare classes controlled by capacity in a risk-averse setting. The revenue management strategy aims at limiting the down-side risk, and in particular, value-at-risk. A value-at-risk optimised policy offers an advantage when considering applications which do not allow for a large number of reiterations. They allow for specifying a confidence level regarding undesired scenarios.

We introduce a computational method for determining policies which optimises the value-at-risk for a given confidence level. This is achieved by computing dynamic programming solutions for a set of target revenue values and combining the solutions in order to attain the requested multi-stage risk-averse policy. We reduce the state space used in the dynamic programming in order to provide a solution which is feasible and has less computational requirements. Numerical examples and comparison with other risk-sensitive approaches are discussed.

Keywords: capacity control, revenue management, risk, value-at-risk

Please reference this paper as:

Matthias Koenig, Joern Meissner. Value-At-Risk Optimal Policies for Revenue Management Problems. Working Paper (available at <http://www.meiss.com>), Lancaster University Management School, 2011.

BIB_T_EX and plain text references are available for download here:
<http://www.meiss.com/en/publications/var-revenue-management.html>

1 Introduction

Revenue management deals with controlling a revenue stream resulting from selling products using a fixed and perishable resource. The industries which use revenue management are manifold. The most popular representatives are airlines, hotels, rental cars, and advertising. But revenue management is also common in event management, ferry lines, retailing or healthcare, to name a few. Talluri and van Ryzin (2005) and Chiang et al. (2007) provide a comprehensive overview of revenue management.

The firm sells multiple products, each consuming a fixed resource with a limited capacity. In this setting, we consider quantity-based revenue management in which a company offers all or just a subset of all products at each point in time. There is a finite time horizon for selling the products, as at the end of the horizon, the salvage value of the resource is zero.

The most common settings use the assumption of a risk-neutral objective. Thus, the policy of the firm is the maximisation of the expected value of its revenue. Often, such a risk-neutral objective is sufficient. As in most applications, such as daily operating ferry lines, this policy is repetitively used. By the law of large numbers, using the expected value as the objective function is then appropriate.

Nevertheless, risk neutrality may not be adequate for other industries, such as event management, that do not support a large number of repetitions of a policy. Several scenarios are known that argue for the considerations of risk-sensitive or risk-averse policies.

Levin et al. (2008) emphasise that, in particular, an event promoter has a high risk, as the promoter cannot count on a large number of reiterations of events. The promoter faces high fixed costs and predominantly has to recover them in order to avoid a possible high loss. Financial and also strategic reasons might not allow running into negative cash, because operational mobility might suffer.

Both Bitran and Caldentey (2003) and Weatherford (2004) provide further examples that risk-neutral considerations are not applied for every real scenario. They report that airline analysts show some natural risk-averse behaviours, and they overrule their revenue management system in situations when the system recommends waiting for high-fare passengers, instead accepting low-fare passengers a few days before flight departure.

That risk-neutral and risk-sensitive policies make a difference is shown in several recent papers. Barz and Waldmann (2007), Huang and Chang (2009), Koenig and Meissner (2009a) and Koenig and Meissner (2009b) analyse both types of policies using the same underlying model that is used in this paper. All four approaches analyse the

effects of applying different kinds of risk-sensitive policies, assuming various levels of risk aversion for a decision maker. However, none of these approaches computes an optimal policy for the common risk measures, such as standard deviation, value-at-risk, or conditional-value-at-risk. However, simulations can be run to determine their values for a given policy.

In this paper, we propose a method which computes a value-at-risk optimal policy. The value-at-risk ($V@R$) is a common risk measure often used in finance (cf. Jorion 2006). It measures down-side risk and is determined for a given probability level. With regard to $V@R$, this probability level is often referred to as confidence level. In our context, the $V@R$ is the lowest revenue which exceeds the confidence level, which is often set at 5 or 10%. Basically, it is a quantile of the revenue distribution determined by the given confidence level.

In order to find a $V@R$ optimal policy, we take advantage of the computation of target level optimal policies as proposed by Koenig and Meissner (2009b). The target level optimal policy can be computed for a certain target and gives information about the probability of not achieving this target. This probability is minimised to find the best policy. It defines a confidence level for a fixed target, which is the corresponding $V@R$. Hence, our task is similar to computing a target level optimised policy, but we optimise the threshold value instead of the percentile. We compute $V@R$ optimal policies and their associated confidence levels. We determine then the policy of the desired confidence level by evaluating the confidence levels of these policies. We describe in this paper how that can be accomplished in an efficient manner.

The advantage of using $V@R$ as a parameter to be optimised is that it is a well-known risk measure, and it is easily interpreted by practitioners. A desired confidence level is specified, and the $V@R$ is returned in the monetary unit of the revenue. Other risk-sensitive approaches often require an interpretation of an uncommon parameter to adjust the desired level of risk preference. $V@R$ is well established and used by risk analysts and decision makers as standard tool not only for financial investments. The risk of a strategy pursued by a decision maker can be assessed by a clear definable risk exposure. This enables risk assessments and planning on an organisational level. Managers can choose their confidence level and communicate it to upper management and investors as well.

Further, a decision maker can define the confidence level to be used for a range of problems although the problems might differ in their settings. This is a great benefit of the $V@R$ approach when compared with the target level approach which might require different target values for each problem setting.

The contribution of this paper is a novel approach in order to assess risk in a revenue management setting. Our approach computes efficiently a value-at-risk optimal policy. To this purpose, we introduce an innovative method in order to reduce the state space of the method which computes a target level optimal policy. We present a simulation study which highlights that our state space reduction still yields high accuracy for the $V@R$ computation even with a significant decreased number of states. In this way, we deliver also a solution which is feasible and has less computational requirements.

The paper is structured as follows. This introduction is followed by a brief overview of related work dealing with revenue management models incorporating risk in Section 2. In Section 3, we continue with the description of the revenue model, which builds our basic position. We describe the target level approach and how we use it to efficiently obtain a $V@R$ optimal policy. We discuss different strategies useful for numerical approximation of such a policy. Section 4 gives a detailed overview of the numerical results and studies the effect of numerical approximation methods. Finally, we conclude this paper in Section 5.

2 Related Work

As a starting point for our analysis we use the basic model by Lee and Hersh (1993). They introduce the dynamic capacity control model in a risk-neutral setting. Lautenbacher and Stidham (1999) take this model further and derive a corresponding Markov decision process. This description as a Markov decision process is advantageous for model extensions.

First risk considerations in revenue management models are proposed by Feng and Xiao (1999). Their model considers risk in terms of variance of sales due to changes of prices. To this end, a penalty function reflecting this variance is incorporated in the objective function of the model. Further, Feng and Xiao (2008) integrate expected utility theory into revenue management models in order to support risk-sensitive decisions.

Expected utility theory as tool for risk consideration is recommended by Weatherford (2004), as well. From a practitioner's perspective, he criticises risk-neutral revenue management, in particular, the expected marginal seat revenue (EMSR) heuristic by Beloba (1989), and endorses risk-averse models.

Barz and Waldmann (2007) base their risk-sensitive model on the Markov decision process of the dynamic capacity model and expected utility theory. They integrate an exponential utility function as the objective function into the Markov decision model. The exponential utility function allows the use of different levels of risk sensitivity.

Barz (2007) points out the use of a utility function with an aspiration level in the same setting but does not discuss the computation of an optimal policy for this utility function. Maximising expected utility using an aspiration level states the same problem as done by the target level objective which is discussed in this paper .

Another way of employing expected utility theory in a revenue management context is proposed by Lim and Shanthikumar (2007). They analyse robust and risk-sensitive control with an exponential utility function for dynamic pricing.

Lai and Ng (2005) formulate a robust optimisation model for revenue management in the hotel industry. Their model incorporates mean versus average deviation. Mitra and Wang (2005) look at mean-variance, mean-standard-deviation and mean-conditional-value-at-risk approach for deriving a risk-sensitive objective function with revenue management application in traffic and networks. Koenig and Meissner (2008) demonstrate that risk considerations might lead to different decisions when deciding between a quantity-based and price-based revenue model.

Also applying risk considerations to the dynamic capacity model, Huang and Chang (2009) show the effect of using a relaxed optimality condition instead of the optimal one. They investigate model behaviour in numerical simulations and discuss results, given as mean and standard deviation and in a ranking based on a Sharpe ratio. A related approach is presented by Koenig and Meissner (2009a), who provide a detailed study of several risk-averse policies for the dynamic capacity model by applying risk measures.

Regarding the use of $V@R$, Lancaster (2003) provides some strong arguments. He demonstrates that risk-neutral revenue management models are vulnerable to the inaccuracy of demand forecasts. Inspired by the $V@R$ metric, he recommends the relative revenue per available seat mile at risk metric. His metric measures the expected maximum of underperformance over a time period for a given confidence level.

Finally, the idea of expanding the state spaces of revenue management models is used by Levin et al. (2008) and Koenig and Meissner (2009b) in order to consider risk in terms of probability for achieving a certain given revenue target. Levin et al. (2008) incorporate risk aversion into a dynamic pricing model of perishable products by integrating constraints into the objective function. Koenig and Meissner (2009a) use the Markov decision model of the dynamic capacity control model and compute optimal policies for revenue targets. Section 3 explains how to find a $V@R$ optimal policy that can employ this model. In a similar manner, finding a $V@R$ optimal policy could also integrate the approach of Levin et al. (2008) for computing the probability of achieving a desired target in the associated context.

3 Modelling and Algorithm

In this section, we begin with a brief introduction of a well-known revenue management problem originally stated in a risk-neutral formulation by Lee and Hersh (1993). We continue with a short summary of a recently proposed modification of this problem which leads to a risk-sensitive model. The risk-sensitive model optimises the risk of failing a previously defined revenue target and provides a basis for the proposed computational approach which focusses on the value-at-risk metric. The value-at-risk metric is explained, and its computation is described in our setting.

3.1 Dynamic Capacity Control Revenue Management Problem

Lee and Hersh (1993) introduce a revenue management model often referred to as the dynamic capacity control model. It was originally formulated for the airline industry, and we also describe it in terms of this industry. Lautenbacher and Stidham (1999) state the problem as a Markov decision process. Using this representation, it is more convenient to derive risk-sensitive policies as done by Barz and Waldmann (2007) for an exponential utility function and by Koenig and Meissner (2009b) for a target level criterion. As we are interested in a computational approach for value-at-risk policies, we focus on dynamic programming equations which can be derived from stating the problem as Markov decision processes.

The model of Lee and Hersh (1993) divides the booking period for a single-leg flight into N decision periods. The decision periods are assumed small enough so that there is no more than one arrival in one period. The decision periods are represented by $n \in \{0, \dots, N\}$ and 0 is the period of departure. There are k different fare classes with fares $F_i, F_1 > F_2 > \dots > F_k$. Further, the probability $p_{n,i}^r$ denotes a request for the fare class i in period n . Probabilities for the last decision period $n = 0$ are zero for all fare classes: $p_{0,i}^r = 0$, meaning the last decision is made at $n = 1$. The probability of no request for any class is given by $p_{n,0}^r = 1 - \sum_{i=1}^k p_{n,i}^r$. Initial seat capacity is C , and remaining seats in time period n are given by $c \leq C$. In this model, a policy π is built from the decision rules which decide to accept or reject a booking request given the current capacity and time. The set of all policies is denoted by Π . The optimal risk-neutral policy $\pi^* \in \Pi$ is the policy which achieves the maximal expected revenue $V_n^{\pi^*}(c) = \max_{\pi} \mathcal{E} \left(\sum_{j=0}^n r_j \right)$, where r_n denotes the random variable for the gained revenue at time n when using a policy π . As Lee and Hersh (1993) show, such an optimal policy can be computed by a dynamic programming solution:

$$V_n^{\pi^*}(c) = \begin{cases} \sum_{i=0}^k p_{n,i}^r \max_{a \in \{0,1\}} \{aF_i + V_{n-1}^{\pi^*}(c-a)\}, & n > 0, c > 0, \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

3.2 Target Level Objective

The risk-sensitive approach proposed by Koenig and Meissner (2009b) builds the basis for calculating a value-at-risk optimised policy. The authors compute an optimal policy for achieving a given target revenue. To this end, they follow a method described by White (1988), Wu and Lin (1999) and Boda and Filar (2006). Boda and Filar (2006) describe the latter approach as a target-percentile problem, as the percentile for a fixed target is optimised.

First, the objective function is the probability of failing the given target revenue. Thus, the objective function has to be minimised in order to derive the risk-sensitive policy. Second, the Markov decision process is augmented by a further state representing the currently remaining target to be achieved in later time steps.

We use the same notation as before and introduce a few more variables. The recent target revenue is denoted by x_n and the given target value to be achieved at N time steps to go is x_N . The value function $W_n^\pi(c, x_n) := \mathcal{P}^\pi \left((\sum_{j=0}^n r_j) \leq x_n \right)$ stands for the probability of failing a target x_n , applying a policy $\pi \in \Pi$ in n remaining time steps and with remaining capacity c . The optimal policy $\tilde{\pi}^* = \operatorname{argmin}_\pi W_N^\pi$ minimises the risk of not attaining the target x_N . The following dynamic programming solution computes this policy:

$$W_0^{\tilde{\pi}^*}(c, x_0) = \begin{cases} 1 & x_0 > 0, \\ 0 & \text{otherwise,} \end{cases}$$

$$W_n^{\tilde{\pi}^*}(c, x_n) = \sum_{i=0}^k p_{n,i}^r \min_{a \in \{0,1\}} \{W_{n-1}^{\tilde{\pi}^*}(c-a, x_n - aF_i)\}. \quad (2)$$

For a target level x_N , we have to consider all possible realisations ending at the final time step 0. With each ongoing time step, a part of the target value can be achieved depending on the decision made. The new target revenue x_{n-1} of the next time step $n-1$ is given by the current target value minus the fare achieved in the current time step $x_n - aF_i$.

The boundary conditions for time step 0 are initialised with 1 for all positive targets and 0 otherwise. For all fares F_i attainable in time step 1, the probability of failing is

less than 1, so their probabilities can be excluded in the sum in Equation 2. Computing $W_N^{\tilde{\pi}^*}$ starts with initialising time step 0 and proceeds to time step N .

Our algorithm calculates the probability of accepting a seat request, which reduces the target by the seat's fare, and the probability of rejecting the request, which retains the current target level. The optimal decision rule either accepts or rejects the request depending on which event has the lower probability.

3.3 Reduction of state space

However, the computation of the dynamic programming solution requires the computation of all cumulative rewards up to the specified target x_N . As this computation of the complete solution is very inconvenient, a more suitable way is using a grid as discussed by Boda et al. (2004). In particular, the state space dimension which represents the target levels is reduced.

To this end, the complete range of all cumulative rewards is discretised. The interval between 0 and the target x_N is separated into m smaller intervals. Each interval spans a width of $\frac{x_N}{m}$. We use $y_i, i \in \{0, \dots, m\}$ as variables for interval boundaries, and the intervals are $[\mathcal{y}_0, \mathcal{y}_1] := [0, \frac{x_N}{m}]$, $[\mathcal{y}_1, \mathcal{y}_2] := [\frac{x_N}{m}, \frac{2x_N}{m}]$, \dots , $[\mathcal{y}_{m-1}, \mathcal{y}_m] := [\frac{x_N(m-1)}{m}, x_N]$. Instead of computing for each possible cumulative reward target x , only the upper boundaries are taken as targets. A target value inside an interval $\mathcal{y} \in (\mathcal{y}_i, \mathcal{y}_{i+1}]$ is rounded to the upper interval boundary \mathcal{y}_{i+1} . This boundary value \mathcal{y}_{i+1} is used while approximately computing the dynamic programming solution.

The computation of $W_n^{\tilde{\pi}^*}$ is done only with value pairs of targets \mathcal{y}_i and probabilities $W_n^{\tilde{\pi}^*}(c, \mathcal{y}_i)$. We obtain a grid of values $\{(\mathcal{y}_0, W_n^{\tilde{\pi}^*}(c, \mathcal{y}_0)), \dots, (\mathcal{y}_m, W_n^{\tilde{\pi}^*}(c, \mathcal{y}_m))\}$ for $c \in \{0, \dots, C\}$. Using dynamic programming Equation 2, the probability values of the grid can be updated in various ways. The simplest method is rounding occurring target values to the upper value, thus $W_n^{\tilde{\pi}^*}(c, \mathcal{y}) = W_n^{\tilde{\pi}^*}(c, \mathcal{y}_{j+1}) \forall \mathcal{y} \in (\mathcal{y}_j, \mathcal{y}_{j+1}]$. However, this is very inaccurate.

We propose using nearest neighbour or linear interpolation as both offer a more accurate way. Nearest neighbour approximation selects the value nearest to the actual required target value \mathcal{y} . If the inequality $|\mathcal{y}_{j+1} - \mathcal{y}| < |\mathcal{y}_j - \mathcal{y}|$ is valid, the upper value on the grid is taken $W_n^{\tilde{\pi}^*}(c, \mathcal{y}) = W_n^{\tilde{\pi}^*}(c, \mathcal{y}_{j+1})$ else the lower value $W_n^{\tilde{\pi}^*}(c, \mathcal{y}) = W_n^{\tilde{\pi}^*}(c, \mathcal{y}_j)$ is taken. Linear interpolation computes weights according to the distances between actual value and grid values. These weights are combined for computing a value for $W_n^{\tilde{\pi}^*}(c, \mathcal{y}) = \frac{|\mathcal{y}_{j+1} - \mathcal{y}|}{\mathcal{y}_{j+1} - \mathcal{y}_j} W_n^{\tilde{\pi}^*}(c, \mathcal{y}_j) + \frac{|\mathcal{y}_j - \mathcal{y}|}{\mathcal{y}_{j+1} - \mathcal{y}_j} W_n^{\tilde{\pi}^*}(c, \mathcal{y}_{j+1})$.

The minimum operator of the dynamic program of Equation 2 is neutral when choosing between an action in case of an equality. There can be several ways for achieving the same minimum, and one of these ways should be selected. If several policies

which achieve the same probability exist, it might be beneficial to choose the policy which achieves the highest revenue of these policies. Regarding the dynamic program, one approach is to accept a request instead of rejecting it when the minimum operator is neutral.

3.4 Value-at-Risk

The target level approach provides us with the means for computing a value-at-risk policy. We explain the value-at-risk metric first and move then to the computation of a value-at-risk optimal policy.

Given a predefined fixed confidence level, the value-at-risk metric computes the maximum loss that one might be exposed to. The confidence level $\alpha \in [0, \dots, 1]$ specifies a probability level and its associated α -quantile is the value-at-risk. There is some inconsistency in the nomenclature of value-at-risk in the literature (cf. Pflug and Römisch 2007a, p57). We use the following definition of the value-at-risk:

$$V@R_\alpha(Y) = \inf\{u : \mathcal{P}(Y \leq u) \geq \alpha\},$$

where Y is a random variable and \mathcal{P} denotes a probability measure. Using this definition, common values for α are 5 or 10 percent.

Applying the $V@R_\alpha$ metric to our model, we use the gained revenue r_n as the random variable and get

$$V@R_\alpha^\pi \left(\sum_{j=0}^n r_j \right) = \inf \left\{ u : \mathcal{P}^\pi \left(\sum_{j=0}^n r_j \leq u \right) \geq \alpha \right\} = \inf \{ u : W_n^\pi(c, u) \geq \alpha \}, \quad (3)$$

with a policy π , remaining time steps n and remaining capacity c .

As we are dealing with revenue, we are interested in finding the policy $\bar{\pi}^*$, which has the maximal $V@R_\alpha$ of all policies Π given confidence α . In other words, we are looking for the policy $\bar{\pi}^*$ which has the highest revenue target of all policies Π given the quantile α . Thus, α fixes the probability of failing a target which has to be determined for every policy $\pi \in \Pi$. The best policy $\bar{\pi}^*$ fails with the same probability α as other policies but achieves a higher target.

The results of Wu and Lin (1999) show that $W_N^{\bar{\pi}^*}(c, x_N)$, as computed by Equation 2, has the property of a cumulative distribution function of variable x_N ; $W_N^{\bar{\pi}^*}(c, x_N)$ is increasing in x_N . Thus, we can employ Equation 2 for computing the policies which optimise the target quantiles of a range of targets. We find the optimal policy for $V@R_\alpha$ by computing a range of best $\bar{\pi}^*$ policies for a range of target values x_N , and

select the policy with the lowest probability $W_N^{\hat{\pi}^*}$ which is yet greater than or equal to α . We can find the $V@R_\alpha$ by using a look-up table or in a similar way by a binary search.

The targets x_N and their associated confidence level $W_N^{\hat{\pi}^*}(c, x_N)$ can be stored in a look-up table. The table is filled by such value pairs, whereby the accuracy of the result depends on the step size used and the interval boundaries used for the various values for x_N . This enables us to look up the target which achieves a quantile equal to confidence α .

Binary search looks up in a sorted sequence for an element by continually splitting the sequence by its median and retaining only the part that the element must be contained in. We can search the $V@R_\alpha$ in a similar way, as $W_N^{\hat{\pi}^*}(c, x_N)$ is an increasing function in x_N . We start with an arbitrary target x_N and decrease or increase it depending on $W_N^{\hat{\pi}^*}(c, x_N)$.

It is also possible to apply binary search to a sequence of targets without precomputation of a full grid. The probabilities of the targets can be computed ‘on the fly’ then. However, the computation of a probability of a target involves the computations of probabilities of targets lower than this target. Therefore, when no precomputation is done caching of already calculated data should be considered in order to avoid repeated computations of the same data.

4 Numerical Results and Discussion

We evaluated the proposed computation method by the same model introduced by Lee and Hersh (1993). Their model serves as an example in various recent papers, cf. Barz and Waldmann (2007), Huang and Chang (2009), Koenig and Meissner (2009a), Koenig and Meissner (2009b). Hence, it provides a basis for a comparison of different policies. Further results are given in the appendix of this paper .

4.1 Experiment Setup

The parameters of this model use $N = 30$ time periods to go before departure. At this point in time, there is a capacity $C = 10$ of seats left. Four fare classes are given with the prices $F_1 = 200, F_2 = 150, F_3 = 120, F_4 = 80$. The probability of request for a distinct fare in the remaining periods are given in Table 1.

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 probabilities of a customer request for fare class i in time period n .

4.2 $V@R$ Computation and Evaluation

We demonstrate a computational approach for finding optimal $V@R$ policies for $\alpha = 5\%$ and $\alpha = 10\%$ as described in the previous section. In this way, we get an achievable $V@R$ value, as well as its corresponding optimal policy.

Table 2 shows the results of computing for a range of possible targets the probabilities of failing them. The underlying computation is based on computing the probability of not achieving a target for every possible target. Thus, no grid which combined ranges of values was used in this case. The first row of Table 2 shows each possible target in the range between 1100 and 1250. This range is just an extract of the overall range of achievable targets. The second row shows the probability of not achieving the target. The next three rows are the simulation results evaluating the policy computed for a target.

We evaluated a policy by using its decision rules in a simulation applying random arrivals according to the probabilities of Table 1. Each simulation result was based on 1000 random runs, and for each set of runs, the same random numbers were used. We used the decision rule of accepting a request if the decision had no effect on the probability. Further, we switched to the risk-neutral policy, if the $V@R$ was attained in a simulation run.

Table 2 shows the fraction of runs which failed the corresponding target, the average and the standard deviation over all achieved revenues. Comparison of the computed probability and the fraction of simulation runs not reaching the target were reasonable within numerical errors.

A possible target represents the $V@R_\alpha$ value and the associated probability, its α value. We find the searched $V@R_{5\%}$ for by looking up the α nearest to 5%, the same way it is done for $\alpha = 10\%$. These determined values-at-risk are highlighted in bold face in Table 2. As the possible targets were not a continuous but a discrete domain, there were also no continuous values for α . Thus, there is no $V@R_{10\%}$ but a $V@R_{10.1\%}$, which is nearest to 10% confidence. This is the same for $\alpha = 5\%$, respectively, but the difference is smaller and not visible in the table.

Target value x_N	1100	1110	1120	1130	1140	1150	1160	1170
$W_N^{\tilde{\pi}^*}$	0.039	0.044	0.047	0.050	0.054	0.060	0.065	0.068
Simulation of target value policies								
Failed target	0.041	0.046	0.048	0.049	0.055	0.059	0.068	0.071
Rev. (average)	1327	1327	1326	1326	1326	1326	1327	1327
Rev.(std. dev.)	166	164	162	161	161	159	158	157
Target value x_N	1180	1190	1200	1210	1220	1230	1240	1250
$W_N^{\tilde{\pi}^*}$	0.074	0.082	0.088	0.093	0.101	0.111	0.120	0.126
Simulation of target value policies								
Failed target	0.072	0.077	0.082	0.089	0.101	0.121	0.125	0.131
Rev. (average)	1328	1329	1329	1330	1331	1333	1335	1336
Rev.(std. dev.)	155	154	153	151	152	152	152	151

Table 2: Extract of the look-up table for finding the $V@R$ nearest to desired values 0.05 and 0.10 for α . Target levels, theoretical percentiles, achieved percentiles (failed target), averages and standard deviation of revenues are shown. The results of the simulations are generated by applying the corresponding policy.

The effect of applying a grid is demonstrated in Table 3. The target level dimension of the state space is reduced by lowering the grid resolution. Results of grid resolutions of $m = 10$, $m = 20$, $m = 40$, $m = 80$ and $m = 166$, which is the highest grid resolution as all possible 166 targets are considered, are compared using exemplarily the confidence value $\alpha = 10\%$. The result of applying the risk-neutral policy is given for comparison. We selected the policy with $W_N^{\tilde{\pi}^*}$, which is nearest and greater than or equal to the desired α . The results in Table 3 show that the inaccuracy increases with decreasing grid resolution. A lower grid resolution results in a lower accuracy of $W_N^{\tilde{\pi}^*}$, and the determined policies π^* do achieve their objective more imprecisely. The standard deviations, which increase with decreasing grid resolution, emphasise this.

Further, the simulation results demonstrated that policies which were computed by linear interpolation with a grid are more suitable for finding a $V@R$ optimal policy for a desired α confidence than policies computed by the nearest neighbour method. Taking into consideration that the state space was strongly reduced, the policies computed by linear interpolation worked quite well with grid sizes down to $m = 20$.

We take a closer look at the different effects of using nearest neighbour selection or linear interpolation in Figures 1 and 2. Both figures show on the axis of abscissae the $V@R_\alpha$ and on the axis of ordinates, the corresponding confidence level α . Each

	$\alpha := W_N^{\tilde{\pi}^*}$ $V@R_\alpha$		Simulation of policy π^*		
			$V@R_{10\%}$	Rev. (avg.)	Rev. (std.)
Risk-neutral	-	-	1130	1408	203
All targets $m = 166$	0.101	1220	1210	1331	152
Linear interpolation					
$m = 80$	0.105	1225	1180	1336	151
$m = 40$	0.126	1250	1140	1346	154
$m = 20$	0.105	1200	1200	1361	157
$m = 10$	0.160	1200	1140	1398	188
Nearest neighbour					
$m = 80$	0.119	1250	1150	1337	152
$m = 40$	0.118	1110	1130	1322	159
$m = 20$	0.199	1200	1070	1309	174
$m = 10$	0.100	1400	1130	1334	162

Table 3: Comparison of approximation methods by using a grid with different resolution and interpolation. Simulation results were generated by applying the determined $V@R$ optimal policy.

depicted graph represents the computed best α for a $V@R_\alpha$ or vice versa. The several graphs in the figures show the effect of using different grid resolutions with nearest neighbour and linear interpolation. Grid resolutions were the same as in Table 3: 166, 80, 40, 20, 10, and for 166, no grid approximation was necessary.

Figure 1 makes it obvious how the accuracy decreased along with decreasing grid resolution when using the nearest neighbour approximation. The graphs of $m = 80$ and $m = 40$ deviated only a little from the accurate graph of $m = 166$. However, the graphs of $m = 20$ and $m = 10$ deviated significantly from the accurate graph and thus, no longer provide reasonable results. This was quite different from the use of linear interpolation.

As expected and shown in Figure 2, linear interpolation provided better approximation results than the nearest neighbour selection. The graph of $m = 80$ nearly matched the graph for accurate resolution, and the graph of $m = 40$ deviated only slightly from it. The first obvious deviation came with the graph $m = 20$ which might be an acceptable approximation. The graph of $m = 10$ deviated strongly and might no longer be a useful approximation in practice. However, linear interpolation was significantly more accurate than nearest neighbour and provided reasonable resolution down to ca. 1/8 of the original and accurate resolution.




Figure 1: Increasing inaccuracy as effect by decreasing grid size when using nearest neighbour.

As linear interpolation was a more accurate approximation than the nearest neighbour selection, we focused on linear interpolation for a further investigation of the impact of grid resolution. Figure 3 displays revenue results from 1000 simulation runs. Using different grid sizes as before, the determined policy for $\alpha = 10\%$ was computed and applied for each simulation run. The axis of the abscissae is the achieved revenue, and the axis of the ordinates is the number of counts the associated revenue was achieved. A histogram shows for a policy of a certain grid resolution the revenue distribution. Further, the results achieved by a risk-neutral policy are given for the purpose of comparison.

Comparing the histograms, we can see that the shape of the revenue distribution of the risk-neutral policy differs from those of the risk-sensitive $V@R_\alpha$ policies for $\alpha = 10\%$. We distinguish between the $V@R_\alpha$ used for finding a policy for $\alpha = 10\%$ and the resulting $V@R_{10\%}$ measurement of the simulation runs. The histograms of the results of the policies of grid resolutions $m = 166$, $m = 80$ and $m = 40$ look very similar in their general shape. We note the peak at revenue of approximately 1250. This was expected as the polices were optimised by ‘moving’ the $V@R_\alpha$ to the






Figure 2: Increasing inaccuracy as effect by decreasing grid size when using linear interpolation.

highest revenue (the right hand side of the distribution) while limiting revenues which are lower the $V@R_\alpha$ (the left hand side of the distribution). However, the policies did not ‘consider’ the shape of the distribution on either side of the $V@R_\alpha$. This resulted in the appearance of the peak near the $V@R_\alpha$.


The results of grid size of $m = 20$ and $m = 10$ were quite interesting. The limited grid resolution seemed no longer possible to ‘shift’ revenues above the $V@R_\alpha$ and the shape of the revenue distributions became similar to that of the risk-neutral policy results. We can see that the histogram of the results from the risk-neutral policy has the largest similarity with that of the results of the policy using $m = 10$. There were differences which yielded consequent different mean revenue and attained $V@R_{10\%}$ of the simulations. The shape of the histogram of the results of the $m = 166$ policy and the shape of the histogram of the risk-neutral policy can be considered as two extremes. By decreasing the grid resolution, the shape of a histogram alters from the one extreme to the other. Thus, the shape of the histogram of the results of the $m = 20$ policy looks like the two extreme shapes merged together.




Risk neutral, $V@R_{10\%}$: 1130




Grid 166, $V@R_{10\%}$: 1210




Grid: 80, $V@R_{10\%}$: 1180



Grid: 40, $V@R_{10\%}$: 1140



Grid: 20, $V@R_{10\%}$: 1200



Grid: 10, $V@R_{10\%}$: 1140

Figure 3: The histograms show the effect of grid resolution on the revenue distribution of numerical simulation. The $V@R_{10\%}$ is given for 1000 simulation runs applying the computed best policy for $\alpha = 10\%$.

However, the achieved $V@R_{10\%}$ of each experiment has to be assessed with the data of Table 3. The policies were the results of an approximation which did not allow every

possible α . The table shows that for $m = 166$, $m = 80$ and $m = 20$ only, the values of $W_N^{\tilde{\pi}^*}$ were 0.101, 0.105 and 0.105, respectively, and thus close to the desired value of $\alpha = 10\%$. Taking this into account, the results of the simulations were consonant with the expected behaviour of the policies.

Hence, only a grid resolution m approximating a policy $\tilde{\pi}^*$ should be chosen which predicts a value $W_N^{\tilde{\pi}^*}$, which has a small difference to the desired confidence level α .

4.3 Comparison with Policies Employing Exponential Utility

Finally, we present a comparison of the $V@R_\alpha$ policies with another risk-sensitive policy, in particular, policies which employ an exponential utility function for implementing risk aversion. We only used this further kind of risk-sensitive policies, because the results of Koenig and Meissner (2009a) show that there are only small differences between the several risk-sensitive policies, including exponential utility based ones. We refer to the paper of Koenig and Meissner (2009a) for a detailed explanation of the other risk-sensitive policies.

A comparison between the presented $V@R_\alpha$ approach and the exponential utility based one is difficult due to their different objectives. Whereas a $V@R_\alpha$ policy maximises the revenue for a certain confidence level α , the exponential utility policy maximises the utility defined by a parameter specified by a level of risk aversion. In order to arrive at comparable experiments, we had to find the level of risk aversion which matches the confidence level. To this end, we ran simulations of utility based policies for a range of levels of risk aversion. For comparison, we chose the utility based policy which achieved the best $V@R_\alpha$ and highest mean revenue. We ran the experiment exemplarily with $\alpha = 10\%$.

Figure 4 shows the difference of the results achieved by both types of policies. The policy which was based on an exponential utility and achieved the highest $V@R_{10\%}$ had a lower $V@R_{10\%}$ than the $V@R_\alpha$ optimised policy. The mean of the results of the utility based policy was higher and its standard deviation was lower than those of the results of the $V@R_\alpha$ based policy. The histograms illustrate these statistical comparisons. The histogram shape of the utility policy that results is broader and more skewed to the right than that of the $V@R_\alpha$ policy results. This shows that the $V@R_\alpha$ policy ‘shifted’ the revenue above the $V@R_\alpha$, whereas the utility policy optimises a utility function and achieved a more balanced distribution.

We show also the conditional-value-at-risk measure ($CV@R_\alpha$) which is the mean of the revenue values lower than the $V@R_\alpha$. The $V@R_\alpha$ policy did not perform better than the utility based one, regarding the $CV@R_\alpha$. This was an important result, as it

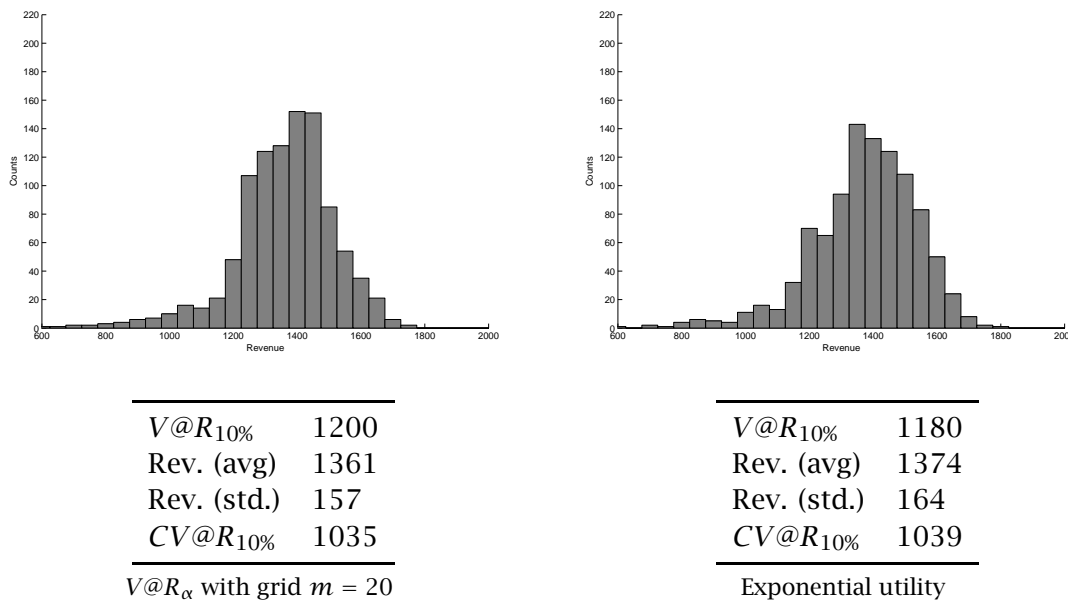


Figure 4: Statistics and histograms of revenue distribution comparing results of $V@R_{\alpha}$ policy with $\alpha = 10\%$ and results of exponential utility based policy with best $V@R_{10\%}$.

illustrated one more time that a $V@R_{\alpha}$ policy does not guarantee the most beneficial lower tail distribution.

5 Conclusions

We have developed a computational approach for finding and approximating the optimal value-at-risk policy for a revenue management problem. The dynamic capacity control model used is one of the quantity-based revenue management models.

Given a confidence level specifying the value-at-risk, the proposed method computes possible value-at-risk results leveraging target level computation and selects the best result fitting the confidence level. In order to reduce computational effort, an approximation method for finding an approximate optimal value-at-risk policy has been proposed.

We have evaluated the proposed approach by computing policies in numerical experiments. A comparison with another risk-sensitive method for the same revenue management problem has been conducted.

The presented methods allow for a fast computation of a good approximation of value-at-risk optimal policies. They provide a basis for applying risk sensitivity in revenue management. However, such policies optimise for value-at-risk but, as often, at

the expense of other measures. This should be borne in mind when applying such policies in practice.

Further, a pure value-at-risk policy cannot suffice in all risk-averse scenarios and a trade-off policy between risk and revenue might be requested. Future work could investigate the computation of such hybrid policies which could be parametrised by confidence level and mean revenue.

Finally, the presented computational approach aiming at value-at-risk optimal policies could also be used for other revenue management models, such as dynamic pricing, if the target level optimal policy is already known.

References

- Barz, C., K.-H. Waldmann. 2007. Risk-sensitive capacity control in revenue management. *Mathematical Methods of Operations Research* **65**(3) 565–579.
- Barz, Christiane. 2007. Risk-averse capacity control in revenue management. Ph.D. thesis, University of Karlsruhe.
- Beloba, P. P. 1989. Application of an probabilistic decision model to airline seat inventory control. *Operations Research* **37**(2) 183–197.
- Bitran, Gabriel, Rene Caldentey. 2003. An overview of pricing models for revenue management. *Manufacturing & Service Operations Management* **5**(3) 203–229.
- Boda, Kang, Jerzy A. Filar. 2006. Time consistent dynamic risk measures. *Mathematical Methods of Operations Research* **63** 169–186.
- Boda, Kang, Jerzy A. Filar, Yuanlie Lin, Lieneke Spanjers. 2004. Stochastic target hitting time and the problem of early retirement. *IEEE Transactions on Automatic Control* **49**(3) 409–419.
- Chiang, W.-C., J.C.H. Chen, X. Xu. 2007. An overview of research on revenue management: Current issues and future research. *International Journal of Revenue Management* **1**(2) 97–128.
- Feng, Y., B. Xiao. 1999. Maximizing revenues of perishable assets with a risk factor. *Operations Research* **47**(2) 337–341.
- Feng, Youyi, Baichun Xiao. 2008. A risk-sensitive model for managing perishable products. *Operations Research* **56**(5) 1305–1311.
- Huang, K., K. Chang. 2009. A simulation analysis for the airline seat control considering expected revenue and risk. *Journal of Revenue and Pricing Management* .
- Jorion, Philippe. 2006. *Value at Risk: The New Benchmark for Managing Financial Risk*. McGraw-Hill.
- Koenig, Matthias, Joern Meissner. 2008. List pricing versus dynamic pricing: Impact on the revenue risk. Tech. rep., Lancaster University Management School.
- Koenig, Matthias, Joern Meissner. 2009a. Risk management policies for dynamic capacity control. Tech. rep., Lancaster University Management School.
- Koenig, Matthias, Joern Meissner. 2009b. Risk minimizing strategies for revenue management problems with target values. Tech. rep., Lancaster University Management School.
- Lai, K.-K., W.-L. Ng. 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.
- Lautenbacher, Conrad J., Shaler Jr. Stidham. 1999. The underlying markov decision process in the single-leg airline yield-management problem. *Transportation Science* 33(2) 136-146.
- Lee, T., M. Hersh. 1993. A model for dynamic airline seat inventory controls with multiple seat bookings. *Transportation Science* 27(3) 252-265.
- Levin, Y., J. McGill, M. Nediak. 2008. Risk in revenue management and dynamic pricing. *Operations Research* 56(2) 326-343.
- Lim, A.E.G., J.G. Shanthikumar. 2007. Relative entropy, exponential utility, and robust dynamic pricing. *Operations Research* 55(2) 198-214.
- Mitra, D., Q. Wang. 2005. Stochastic traffic engineering for demand uncertainty and risk-aware network revenue management. *IEEE/ACM Transactions on Networking* 13(2) 221-233.
- Pflug, Georg Ch., Werner Römisch. 2007a. *Modeling, Measuring and Managing Risk*. World Scientific.
- Talluri, K.T., G.J. van Ryzin. 2005. *The Theory and Practice of Revenue Management*. Springer.
- Weatherford, L.R. 2004. EMSR versus EMSU: Revenue or utility? *Journal of Revenue and Pricing Management* 3(3) 277-284.
- White, D. J. 1988. Mean, variance, and probabilistic criteria in finite markov decision processes: A review. *Journal of Optimization Theory and Applications* 56(1) 1-29.
- Wu, Congbin, Yuanlie Lin. 1999. Minimizing risk models in markov decision processes with policies depending on target values. *Journal of Mathematical Analysis and Applications* 231 47-67.

Appendix

We conducted further experiments to evaluate the effect of a reduction of the state space. To this purpose, we ran experiments with $N = 30$ time periods and $C = 10$ initial capacity. The time periods were divided into five parts of six successive periods of the same request probabilities. We used three different scenarios of different fare structures, see Table 4. The request probabilities of an experiment were generated randomly and that experiment categorised according its load factor $\rho = \frac{1}{C} \sum_{n=0}^N \sum_{j=1}^k p_{jn}$. Note the experiment of Section 4 had a load factor $\rho = 1.32$.

Table 4 summarises the results of the experiments. We accepted confidence levels of $\alpha \in [9.5\%, 10.5\%]$ as not every confidence level was possible in an experiment. The columns of $\alpha := W_N^{\hat{\pi}^*}$ represent averages of 100 theoretically achievable confidence levels of experiments which had a load factor ρ in the range as shown for each row. The columns of $\bar{\alpha}$ show the confidence levels which were achieved by the numerical experiments which used the $V@R_\alpha$ corresponding to that α . Results of an experiment were averaged over 1000 sample runs. In order to obtain a value of $\bar{\alpha}$, 100 experiments were averaged for each category of a load factor, too. For instance, the $\bar{\alpha}$ value 0.101 given in the first row and third column of scenario S1 is the average of 100 different experiments which had a load factor between $[0.75, 1.00)$, where each experiment was averaged over 1000 sample runs.

Further, Table 4 shows the results of the same experiments but with different underlying grid sizes. We used grid sizes without any state space reduction, thus, with enough states for every possible revenue as references. The reduction of those grid sizes is given by the same experiments which used only 50% and 25% of those full grid sizes, respectively. The $V@R_\alpha$ was computed using the full grid size.

We observed that there were only minimal differences between the theoretically achievable α values and the obtained $\bar{\alpha}$ values when the full grid sizes were used (less than 0.5 percent points) independent of the scenario and the load factor. We noticed that the differences between the predicted α values and the obtained $\bar{\alpha}$ values increased with decreasing grid size although there was no clear relationship between the load factors and those differences. The maximal difference was 2 percent points for 50% of full grid size, and 3.2 percent points for 25% of full grid size, respectively.

The deviations of the α values of the different grid sizes were caused due to the fact that the computations of the $V@R_\alpha$ values were done using the full grid size. Those computed $V@R_\alpha$ values were then used for the lower grid sizes as well.

$\rho \in$	full grid		50% of full grid		25% of full grid	
	$\alpha := W_N^{\hat{\pi}^*}$	$\bar{\alpha}$	$\alpha := W_N^{\hat{\pi}^*}$	$\bar{\alpha}$	$\alpha := W_N^{\hat{\pi}^*}$	$\bar{\alpha}$
Scenario S1 : $F = (200, 150, 100)$						
[0.75, 1.00)	0.102	0.101	0.138	0.146	0.154	0.158
[1.00, 1.25)	0.102	0.100	0.133	0.133	0.156	0.128
[1.25, 1.50)	0.103	0.103	0.153	0.144	0.169	0.137
[1.50, 1.75)	0.102	0.100	0.161	0.141	0.174	0.149
Scenario S2 : $F = (200, 150, 120, 80)$						
[0.75, 1.00)	0.102	0.104	0.107	0.111	0.115	0.125
[1.00, 1.25)	0.103	0.105	0.108	0.114	0.118	0.127
[1.25, 1.50)	0.103	0.106	0.111	0.120	0.122	0.133
[1.50, 1.75)	0.102	0.106	0.114	0.124	0.125	0.142
Scenario S3 : $F = (240, 200, 160, 120, 80)$						
[0.75, 1.00)	0.102	0.102	0.128	0.135	0.151	0.159
[1.00, 1.25)	0.103	0.102	0.123	0.126	0.141	0.132
[1.25, 1.50)	0.103	0.102	0.128	0.128	0.149	0.133
[1.50, 1.75)	0.103	0.103	0.133	0.131	0.162	0.137

Table 4: Comparison how grid size reduction affects accuracy of $V@R$ in achieving confidence level α in different scenarios: The values of $\alpha := W_N^{\hat{\pi}^*}$ are averaged theoretically achievable results and $\bar{\alpha}$ are averaged results from numerical experiments. Full grid size means no reduction of states and 50%, 25% of full grid size represents the number of states related to the full grid size, respectively.